UNIVERSIDAD DE CHILE
FACULTAD DE CIENCIAS FÍSICAS Y MATEMÁTICAS
DEPARTAMENTO DE INGENIERÍA MATEMÁTICA

LOCAL DYNAMICS FOR THE CALCULATION OF GLOBAL PROPERTIES IN GRAPHS

TESIS PARA OPTAR AL GRADO DE
DOCTORA EN CIENCIAS DE LA INGENIERÍA,
MENCIÓN MODELACIÓN MATEMÁTICA

LAURA MAYELY LEAL CHACÓN

PROFESOR GUÍA:
IVÁN RAPAPORT ZIMERMANN

PROFESOR CO-GUÍA:
PEDRO MONTEALEGRE

MIEMBROS DE LA COMISIÓN:
AXEL OSSES ALVARADO
MARTÍN RÍOS-WILSON

Este trabajo ha sido financiado por Beca Doctorado Nacional ANID 2019-21191440 y CMM ANID BASAL FB210005.

SANTIAGO DE CHILE
2022

RESUMEN DE LA TESIS PARA OPTAR AL GRADO DE
DOCTORA EN CIENCIAS DE LA INGENIERÍA,
MENCIÓN MODELACIÓN MATEMÁTICA
POR LAURA MAYELY LEAL CHACÓN
FECHA: 2022
PROF. GUÍAS: IVÁN RAPAPORT, PEDRO MONTEALEGRE

DINÁMICAS LOCALES PARA EL CÁLCULO DE PROPIEDADES GLOBALES EN GRAFOS

El problema de clasificación de densidad en grafos consiste en encontrar una dinámica local tal que, dado un grafo y una configuración inicial de 0's y 1's asignada a los nodos del grafo, la dinámica converja a la configuración de punto fijo 1's si la fracción de 1 es mayor que una densidad crítica (típicamente 1/2) y, de lo contrario, converja a la configuración de punto fijo 0's. Para resolver este problema seguimos la idea propuesta en un trabajo anterior [7], donde los autores diseñaron un autómata celular inspirado en dos mecanismos: difusión y amplificación. Aplicamos este enfoque a diferentes clases de grafos bien conocidos, en los que se encuentran grafos completos, regulares, de estrella, Erdös-Rényi y Barabási-Albert.

Un conjunto independiente maximal (MIS) es un conjunto maximal por inclusión de vértices no adyacentes por pares. El cálculo de un MIS es uno de los problemas centrales de la computación distribuida. En esta tesis presentamos y analizamos un algoritmo aleatorio distribuido de estado finito para calcular un MIS en grafos arbitrarios no dirigidos. Nuestro algoritmo es autoestabilizante, es decir, alcanza una salida correcta en cualquier configuración inicial. Analizamos el tiempo de convergencia del algoritmo propuesto, mostrando que en muchos casos el algoritmo converge en tiempo logarítmico con alta probabilidad.

SUMMARY OF THE THESIS TO OBTAIN THE DEGREE OF
DOCTORA EN CIENCIAS DE LA INGENIERÍA,
MENCIÓN MODELACIÓN MATEMÁTICA
BY LAURA MAYELY LEAL CHACÓN
DATE: 2022
ADVISORS: IVÁN RAPAPORT, PEDRO MONTEALEGRE

LOCAL DYNAMICS FOR THE CALCULATION OF GLOBAL PROPERTIES IN GRAPHS

The density classification problem on graphs consists in finding a local dynamics such that, given a graph and an initial configuration of 0's and 1's assigned to the nodes of the graph, the dynamics converge to the fixed point configuration of all 1's if the fraction of 1's is greater than a critical density (typically 1/2) and, otherwise, it converges to the all 0's fixed point configuration. To solve this problem we follow the idea proposed in a previous work [7], where the authors designed a cellular automaton inspired by two mechanisms: diffusion and amplification. We apply this approach to different well-known graph classes, including complete, regular, star, Erdös-Rényi and Barabási-Albert graphs.

A Maximal Independent Set (MIS) is an inclusion maximal set of pairwise non-adjacent vertices. The computation of an MIS is one of the core problems in distributed computing. In this thesis, we introduce and analyze a finite-state distributed randomized algorithm for computing a MIS on arbitrary undirected graphs. Our algorithm is self-stabilizing, that is, reaches a correct output on any initial configuration and can be implemented on systems with very scarce conditions. We analyze the convergence time of the proposed algorithm, showing that in many cases the algorithm converges in logarithmic time with high probability.

A mi amado esposo por el apoyo y contención.

Agradecimientos

Quiero comenzar agradeciendo a Dios por permitirme llegar hasta este punto de mi vida académica.

A mi amado esposo por brindarme su incondicional apoyo.

Al Programa de Doctorado en Ciencias de la Ingeniería, Mención Modelación Matemática de la Universidad de Chile por su financiamiento en matrícula y a su cuerpo docente por el tiempo y conocimiento brindado. Además, agradecer al gobierno de Chile por su financiamiento a través de la beca ANID 2019-21191440.

A los profesores Iván Rapaport y Pedro Montealegre, por el apoyo a lo largo de todo el trabajo de tesis.

Un agradecimiento especial al profesor Axel Osses por la colaboración con el trabajo de investigación de la presente tesis y al equipo encargado de la infraestructura del supercómputo del NLHPC (ECM-02) por permitirme usar sus instalaciones para ejecutar mis programas.

Finalmente y no menos importante gracias a todos aquellos que de una u otra manera, sabiéndolo o no, ayudaron a la realización de este trabajo el cual culminó con dos trabajos publicados [31] y [24].

Contents

Chapter 1

Introduction

1.1 Density Classification

One of the simplest, regular graph topology is a torus. A cellular automaton (CA) is nothing but a local dynamics applied to a torus (or to an infinite grid) [39, 37]. More precisely, a CA consists of n^d cells arranged uniformly spaced in the d-dimensional torus and following a local rule identical in every cell. This local rule, which specifies how the state of each cell is updated as a function of the states of its neighbor cells, is applied in parallel and in discrete time steps.

In the *density classification problem*, the challenge is to find a CA such that, given any initial configuration x^0 of 0's and 1's, it converges to the all 1's fixed point configuration if the fraction of 1's in x^0 is greater than ρ_c and it converges the the all 0's fixed point configuration otherwise. The number $0 < \rho_c < 1$ denotes the *critical density*.

The problem was first formulated for dimension $d = 1$ (a ring) and critical density $\rho_c = 1/2$ [36]. The best-known two-state CA for tackling this instance gave a good solution which was not perfect [20, 14]. In fact, an impossibility result which says that there is no perfect density classifier with two states was later obtained [30].

The impossibility of finding perfect classifiers led many researchers to use different evolutionary computation approaches to evolve good *approximate solutions* [26, 35, 13, 38, 41]. But in order to obtain perfect density classifiers, researchers were forced to modify the original problem [9, 15, 40]. Another idea was to allow the existence of more than one local rule [19, 34]. Fatès designed a two-state *stochastic*CA that solves the density classification problem with arbitrary precision [18].

In [7] the authors used a continuous approach for solving *deterministically* the density classification problem. More precisely, the idea was to use *local averaging* and *saturation*, a process represented by a bistable heat equation. This bistable model, which exhibits two stable critical points (0 and 1), is a particular case of a reaction-diffusion equation widely used for studying phase transitions and front propagation [8, 28, 42]. Similar approaches have been previously used [11, 16]. We refer to [12] for a comprehensive survey about the

density classification problem.

The idea of the present work is to adapt this approach, which was created to the torus topology, to other topologies corresponding to well-known graph classes: complete, regular, star, Erdös-Rényi and Barabási-Albert graphs.

We denote the local dynamics adapted to these graph classes by Φ. Note that Φ corresponds to the discretization of the heat equation and it can in fact be applied to *any* graph.

The idea is the following. Given a critical density $\rho_c \in (0, 1)$, we build a discrete dynamics Φ over an arbitrary connected graph G based on a discrete version of the following equation:

$$\frac{\partial u}{\partial t} - \nu \Delta u = \gamma b_{\rho_c}(u), \tag{1.1}$$

where $u(x, t)$ is the state at time $t \geq 0$ at point x in a domain Ω. The parameter $\nu > 0$ is a diffusion coefficient, $\gamma > 0$ is an amplification parameter and b_{ρ_c} is some suitable bistable function. In this paper we choose the cubic polynomial:

$$b_{\rho_c}(u) = u(1 - u)(u - \rho_c).$$

As for the case of the torus, the resulting nonlinear heat equation exhibits two stable critical points (0 and 1, attractors) and one unstable critical point $(\rho_c, repulsor)$.

The main theoretical result of this thesis says that, for every connected graph G, there are parameters for the large diffusion and small amplification dynamics Φ, such that Φ solves perfectly the density classification problem in G.

With respect to experimental results, we show how the topology of a given graph G influences the convergence time. First of all, the study of the dynamics *without* amplification -which does not solve the density classification problem-indicates that the convergence time depends on the edge density of G, and that the behavior is fundamentally different on trees (in our simulations, the case of trees is provided by star graphs and a subclass of Barabási-Albert graphs).

Then, we give experimental results when a small amplification is indeed present in Φ. We observe a dramatic difference on the convergence time for the classes that are trees with respect to those that are not (the classes that are not trees, in terms of convergence time, are almost indistinguishably). Finally, we study the influence of the amplification factor on the effectiveness and the convergence time of the local dynamics. Again, we see a dichotomy between trees and graphs with cycles.

1.2 Maximal Independent Set

The Maximal Independent Set problem (MIS) is one of the main problems in distributed computing. In its simplest version, it consists in finding an inclusion-maximal set of pairwise

non-adjacent vertices on an undirected graph. The MIS corresponds to a specific case of a wide problem in distributed graph algorithms, known as symmetry breaking. When a distributed algorithm is executed, the nodes of a distributed system are assumed to be in the same state, but in the successive time-steps the nodes are expected to play different roles, hence breaking the symmetry.

The MIS has a trivial solution in the classical sequential setting, where a greedy algorithm sequentially picks an arbitrary vertex, includes it in the maximal independent set, and removes that vertex together with all its neighbors. In the 80's, Karp and Wigderson [27] mentioned that the MIS is an interesting problem in non-centralized computation. Soon after that, Luby [33] and Alon, Babai, and Itai [3] presented simple distributed randomized algorithms solving MIS in $\mathcal{O}(\log n)$ time. Since then, this problem has been studied extensively in the distributed setting. In the LOCAL model, the fastest deterministic MIS algorithms for general graphs run in $\mathcal{O}(\log^5 n)$ [22], and $\mathcal{O}(\Delta + \log^* n)$ time [4]. Ghaffari [21] also obtained a $\mathcal{O}(\log \Delta) + 2^{\mathcal{O}(\sqrt{\log \log n})}$ time randomized algorithm on general graphs, and a $\mathcal{O}(\log a + \sqrt{\log n})$ time randomized algorithm for graphs of arboricity a. With respect to lower bounds, Linial [32] proved that computing an MIS on an n-cycle requires time $\Omega(\log^* n)$. Moreover, Kuhn, Moscibroda and Wattenhofer [29] showed a $\Omega(\sqrt{\log n})$ lower-bound on the round complexity on general graphs.

Another branch of research regarding the MIS problem in the distributed setting consists in considering models with limited resources. One example is the *beeping model* [1], where the nodes are limited to an extremely harsh system of communication. On each round, a node can either broadcast a signal (a beep) or hear whether a neighbor emitted a beeping signal, but noes it is not capable of distinguishing the number nor the sources of the beeping signals it receives. In this model, Afek et al. [1] showed that an MIS can be computed in time $\mathcal{O}(\log^3 n)$ when the nodes have to know the size of the graph n and have $\text{poly}(\log n)$ sized memory. The *stone-age model* is another relevant model with limited resources, where the memory of each node is limited to a constant not depending on the size of the graph. In this model, Emek and Wattenhofer [17] give an MIS algorithm with a running time of $\mathcal{O}(\log^2 n)$. Interestingly, this algorithm requires that the nodes start in a particular initial state in order to be capable of performing correct computation.

An algorithm is called *self-stabilizing* [25] if it can reach a correct output starting from any initial state. The motivation for this kind of algorithm is the capacity of distributed systems to self-repair a faulty configuration when one of the parties crashes. For instance, consider a maximal independent set, where one of the nodes crashes. It is possible that the rest of the nodes do not form a maximal independent set in the remaining graph (for instance, consider an MIS in a complete graph where the unique marked node crashes). In that context, a self-stabilizing MIS algorithm should be able to reach an MIS for any initial state configuration of the nodes.

In a keynote talk of SIROCCO 2022 [23], George Giakkoupis presented two extremely simple randomized algorithms for MIS. His algorithms have two interesting properties: they are self-stabilizing and they require only two or three states. Despite this simplicity, the algorithm has hardly been studied before, and its convergence time is not been settled yet. In this article, we propose a variant of the algorithm of [23] and study its convergence time both numerically and analytically.

Chapter 2

A Large Diffusion and Small Amplification Dynamics for Density Classification on Graphs

The density classification problem on graphs consists in finding a local dynamics such that, given a graph and an initial configuration of 0's and 1's assigned to the nodes of the graph, the dynamics converge to the fixed point configuration of all 1's if the fraction of 1's is greater than a critical density (typically 1/2) and, otherwise, it converges to the all 0's fixed point configuration. To solve this problem we follow the idea proposed in [7], where the authors designed a cellular automaton inspired by two mechanisms: diffusion and amplification. We apply this approach to different well-known graph classes: complete, regular, star, Erdös-Rényi and Barabási-Albert graphs.

2.1 Introduction

One of the simplest, regular graph topology is a torus. A cellular automaton (CA) is nothing but a local dynamics applied to a torus (or to an infinite grid) [39, 37]. More precisely, a CA consists of n^d cells arranged uniformly spaced in the d-dimensional torus and following a local rule identical in every cell. This local rule, which specifies how the state of each cell is updated as a function of the states of its neighbor cells, is applied in parallel and in discrete time steps.

In the *density classification problem*, the challenge is to find a CA such that, given any initial configuration x^0 of 0's and 1's, it converges to the all 1's fixed point configuration if the fraction of 1's in x^0 is greater than ρ_c and it converges the the all 0's fixed point configuration otherwise. The number $0 < \rho_c < 1$ denotes the *critical density*.

The problem was first formulated for dimension $d = 1$ (a ring) and critical density $\rho_c = 1/2$ [36]. The best-known two-state CA for tackling this instance gave a good solution which was not perfect [20, 14]. In fact, an impossibility result which says that there is no perfect density

classifier with two states was later obtained [30].

The impossibility of finding perfect classifiers led many researchers to use different evolutionary computation approaches to evolve good *approximate solutions* [26, 35, 13, 38, 41]. But in order to obtain perfect density classifiers, researchers were forced to modify the original problem [9, 15, 40]. Another idea was to allow the existence of more than one local rule [19, 34]. Fatès designed a two-state *stochastic* CA that solves the density classification problem with arbitrary precision [18].

In [7] the authors used a continuous approach for solving *deterministically* the density classification problem. More precisely, the idea was to use *local averaging* and *saturation*, a process represented by a bistable heat equation. This bistable model, which exhibits two stable critical points (0 and 1), is a particular case of a reaction-diffusion equation widely used for studying phase transitions and front propagation [8, 28, 42]. Similar approaches have been previously used [11, 16]. We refer to [12] for a comprehensive survey about the density classification problem.

The idea of the present work is to adapt this approach, which was created to the torus topology, to other topologies corresponding to well-known graph classes: complete, regular, star, Erdös-Rényi and Barabási-Albert graphs.

We denote the local dynamics adapted to these graph classes by Φ. Note that Φ corresponds to the discretization of the heat equation and it can in fact be applied to *any* graph.

The idea is the following. Given a critical density $\rho_c \in (0, 1)$, we build a discrete dynamics Φ over an arbitrary connected graph G based on a discrete version of the following equation:

$$\frac{\partial u}{\partial t} - \nu \Delta u = \gamma b_{\rho_c}(u), \tag{2.1}$$

where $u(x, t)$ is the state at time $t \geq 0$ at point x in a domain Ω. The parameter $\nu > 0$ is a diffusion coefficient, $\gamma > 0$ is an amplification parameter and b_{ρ_c} is some suitable bistable function. In this paper we choose the cubic polynomial:

$$b_{\rho_c}(u) = u(1 - u)(u - \rho_c).$$

As for the case of the torus, the resulting nonlinear heat equation exhibits two stable critical points (0 and 1, attractors) and one unstable critical point $(\rho_c, repulsor)$.

The main theoretical result of this thesis says that, for every connected graph G, there are parameters for the large diffusion and small amplification dynamics Φ, such that Φ solves perfectly the density classification problem in G.

With respect to experimental results, we show how the topology of a given graph G influences the convergence time. First of all, the study of the dynamics *without* amplification -which does not solve the density classification problem-indicates that the convergence time depends on the edge density of G, and that the behavior is fundamentally different on trees (in our simulations, the case of trees is provided by star graphs and a subclass of Barabási-Albert graphs).

Then, we give experimental results when a small amplification is indeed present in Φ. We observe a dramatic difference on the convergence time for the classes that are trees with respect to those that are not (the classes that are not trees, in terms of convergence time, are almost indistinguishably). Finally, we study the influence of the amplification factor on the effectiveness and the convergence time of the local dynamics. Again, we see a dichotomy between trees and graphs with cycles.

2.2 Preliminaries

In this section we give the main definitions of the concepts and problems used in the rest of the paper.

We consider only finite simple undirected graphs $G = (V, E)$, where V is the set of vertices and $E \subseteq \binom{V}{2}$ is the set of (undirected) edges. For a node v in a graph $G = (V, E)$, we call $N(v) = \{u \in V : \{u, v\} \in E\}$ the neighborhood of v. The cardinality of $N(v)$ is called the *degree* of v and it is denoted $d(v)$. The *maximum degree* is denoted $\Delta(G)$ and corresponds to $\max_{v \in V} d(v)$. Given a graph G with $V(G) = \{v_1, \ldots, v_n\}$, the *adjacency matrix* A of G corresponds to a square matrix of order n such that $A_{ij} = 1$ when $\{v_i, v_j\} \in E(G)$, and $A_{ij} = 0$ otherwise.

A sequence of vertices $P = v_1, \ldots, v_k$ is called a v_1, v_k- *path* if $\{v_i, v_{i+1}\}$ is an edge of G, for each $1 \leq i < k$. Two vertices u and v are *connected* if there exists a u, v-path in G. Being connected defines an *equivalence relation*, and the equivalence classes of this relation (i.e. inclusion maximal sets of connected vertices) are called *connected components* of G. In this article, we consider only connected graphs. A *cycle* is a sequence of $k \geq 3$ vertices $v_1, \ldots, v_k$ that form a v_1, v_k-path and where $\{v_1, v_k\}$ is an edge of G. A connected graph without cycles is called a *tree*.

We denote by $[n]$ the set $\{1, \ldots, n\}$. Given a graph $G = ([n], E)$, a *configuration* is a vector $u \in [0, 1]^n$. In this paper, we study time-discrete dynamical systems over the configurations of graphs. Formally, we study the dynamics given by certain *local updating rules*, which are functions $F : [0, 1]^n \to [0, 1]^n$. Such a function F defines a sequence of configurations given by:

$$u^t = F(u^{t-1}),$$

for each $t \geq 1$, where $u^0 \in [0, 1]^n$ is denoted the *initial configuration*.

Given a configuration $u \in [0, 1]^n$, the *density of* u, denoted by $\rho(u)$, is the quantiy

$$\rho(u) = \sum_{i \in [n]} \frac{u_i}{n}.$$

In that context, given a *critical density* $\rho_c \in (0, 1)$, the *density classification problem* consists in the design of a time-discrete dynamical system F such that, for every Boolean initial configuration $u \in \{0, 1\}^n$:

- If $\rho(u) < \rho_c$ then $\lim_{t \to \infty} F^t = [0, \ldots, 0]$

- If $\rho(u) > \rho_c$ then $\lim_{t \to \infty} F^t = [1, \ldots, 1]$

In other words, a rule F solves the density classification problem if, for every Boolean initial configuration, the dynamics converge to either an all 1's or all 0's configuration, depending on whether the density of the initial configuration is below or above a given critical density.

We also consider an approximated version of the density classification problem. Given $\varepsilon > 0$, we say a rule F solves the ε-approximation of the density classification problem if

- If $\rho(u) < \rho_c - \varepsilon$ then $\lim_{t \to \infty} F^t = [0, \ldots, 0]$
- If $\rho(u) > \rho_c + \varepsilon$ then $\lim_{t \to \infty} F^t = [1, \ldots, 1]$

2.3 A Local Dynamics for Density Classification

In this section we define the large diffusion and small amplification dynamics. They are defined by a local rule Φ is based on the discretization of a bistable nonlinear heat equation. More precisely, given a critical density $\rho_c \in (0, 1)$, the idea is to build a discrete dynamic over an arbitrary connected graph G based on a discrete version of the following equation:

$$\frac{\partial u}{\partial t} - \nu \Delta u = \gamma b_{\rho_c}(u), \tag{2.2}$$

where $u(x, t)$ is the state at time $t \geq 0$ at point x in a domain Ω. The parameter $\nu > 0$ is a diffusion coefficient, $\gamma > 0$ is an amplification parameter and b_{ρ_c} is some suitable bistable function. In this paper we choose the cubic polynomial:

$$b_{\rho_c}(u) = u(1 - u)(u - \rho_c)$$

The resulting nonlinear heat-equation is called the bistable heat equation, since it exhibits two stable critical points (0 and 1, attractors) and one unstable critical point $(\rho_c, repulsor)$.

Let $G = (V, E)$ be a graph with vertex set $V = \{1, \ldots, n\}$ and edge set $E \subseteq \binom{V}{2}$. We denote by $N(i)$ the set of neighbors of node i, and by d_i the degree (number of neighbors) of i. We can discretize Eq. (2.2) with an explicit finite differences scheme on a uniform space difference $h > 0$ and discrete time steps $t_k = k\Delta t$ for some $\Delta t > 0$, obtaining

$$\frac{u_i^{k+1} - u_i^k}{\Delta t} + \frac{\nu}{h^2} \sum_{j \in N(i)} L_{ij} u_j^k = \gamma b_{\rho_c}(u_i^k),$$

where $L = L(G)$ is the Laplacian matrix of G, which corresponds to an $n \times n$ matrix is defined by,

$$L_{ij} := \begin{cases} d(i) & \text{si } \ i = j, \\ -1 & \text{si } \ ij \in E, \\ 0 & \text{si } \ ij \notin E. \end{cases}$$

If we denote by $\sigma = \Delta t\gamma$, $\sigma' = {}^{\nu\Delta t}/_{h^2}$ and $b_{\rho_c}(u_i) = u_i(1-u_i)(u_i-\rho_c)$, we obtain the local rule

$$u_i^{k+1} = u_i^k - \sigma' \sum_{j\in N(i)} L_{ij}u_j^k + \sigma b_{\rho_c}(u_i^k).$$

Finally we choose the parameter ν in order to fix $\sigma' = {}^1/_{(\Delta+1)}$, where $\Delta = \Delta(G)$ is the maximum degree of G. Then, the *large diffusion and small amplification dynamic over* G updating rule Φ_σ:

$$v^t = \left(I - \frac{L}{\Delta+1}\right) u^t, \tag{2.3}$$

$$u_i^{t+1} = f_\sigma(v_i^t), \tag{2.4}$$

where $f_\sigma(x) = x + \sigma b_{\rho_c}(x)$.

Theorem 2.1 *For every connected graph G and for every $\varepsilon > 0$ there exists a $\sigma > 0$ such that the large diffusion and small amplification dynamic over G solves the ε-approximation of the density classification problem.*

The proof of Theorem 2.1 is analogous to the proofs given in [7]. For sake of completeness we include all the details. In order to proof Theorem 2.1, we prove first some technical lemmas. In the first lemma, we show that the dynamic without amplification (i.e. when $\sigma = 0$) converges to a configuration where the state of every vertex is the density of the initial configuration.

In the following, we denote by C the matrix $\left(I - \frac{L}{\Delta+1}\right)$.

Lemma 2.2 *For every $u \in \{0,1\}^n$, $\lim\limits_{t\to\infty} \Phi_0^t(u) = \rho(u)[1,\ldots,1]^T$,*

Proof. When we consider the dynamics without any amplification $\sigma = 0$, we obtain that for every $t > 0$, theupdating rule becomes:

$$u^t = \Phi_0(u^t) = Cu^{t-1} = C^t u^0$$

Observe that C is a doubly stochastic (its entries are not non-negative and the sum of their rows and columns is 1) and symmetric matrix. Moreover, C is primitive, meaning that there exists an $m > 0$ such that $(C^m)_{ij} \neq 0$ for all i, j. The existence of such m simply follows from the fact that the underlying graph is connected. Because of C is symmetric, we know that all of its eigenvalues are real.

From the properties of C and the Perron-Frobenius Theorem we deduce the following properties. First, $\lambda = 1$ is an eigenvalue of C of multiplicity one. Also, the eigenspace associated with the eigenvalue $\lambda = 1$ is spanned by $[1\cdots 1]^T$. The absolute value of all the other eigenvalues is strictly less than 1. We obtain that C can be decomposed as $M^T DM$, where D is diagonal, the elements on the main diagonal of D are the eigenvalues of C and

M is orthonormal. Moreover,

$$\lim_{t\to\infty} D^t = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix}$$

implying that

$$\lim_{t\to\infty} C^t = \frac{U}{n},$$

where U is the square matrix of order n where all its entries are ones. We deduce that for every initial configuration $u^0 \in \{0,1\}^n$,

$$\lim_{t\to\infty} u^{t+1} = \lim_{t\to\infty} C^t u^0 = \frac{1}{n}\left(\sum_{i=1}^{n} u_i^0\right)[1\cdots 1]^T = \rho(u^0)[1\cdots 1]^T.$$

Implying that when $\sigma = 0$ we have that the dynamics converge to a configuration where the state of every node corresponds to the density $\rho(u^0)$ of the initial configuration u^0 □

Let us now study the case when $\sigma > 0$. When $\sigma > 0$ we have that f_σ is strictly increasing, $f_\sigma(0) = 0$ and $f_\sigma(1) = 1$. Therefore, by continuity f_σ is a one-to-one map from $[0,1]$ to $[0,1]$. In the next lemma, we show that when the initial configuration has each coordinate greater than (resp. smaller than) the critical density ρ_c, then the dynamic correctly solves the density classification problem.

Lemma 2.3 *Let $\sigma > 0$. If $u \in [0,1]^n$ is such that, for every $i \in [n]$, $u_i > \rho_c$ then,*

$$\lim_{t\to\infty} \Phi_\sigma^t(u) = [1,\ldots,1]^n$$

Similarly, if $u \in [0,1]^n$ is such that, for every $i \in [n]$, $u_i < \rho_c$ then,

$$\lim_{t\to\infty} \Phi_\sigma^t(u) = [0,\ldots,0]^n$$

Proof. Let us consider first a uniform initial configuration $u^0 = [q,\ldots,q]^T$ with $q \in (\rho_c, 1]$. Then $v^0 = Cu^0 = [q,\ldots,q]^T = u^0$. Since f_σ is increasing, we obtain that for every $i \in [n]$, $u_i^0 < f_\sigma(u_i^0) = \Phi_\sigma(u)_i \leq 1$. Inductively, for every $t \geq 1$ and every $i \in [n]$, $\Phi_\sigma^{t-1}(u)_i < \Phi_\sigma^t(u)_i \leq 1$. We deduce that $\lim_{t\to\infty} \Phi_\sigma^t(u) = [1,\ldots,1]^T$.

Now let us fix a configuration $u \in [0,1]^n$ such that $\rho_c < u_i$ for all $i \in [n]$, and let us call $q = \min_{i\in[n]} u_i \in (\rho_c, 1]$. Then $v = Cu$ satisfies that $\rho_c < q \leq v_i$ for every $i \in [n]$. Therefore $\rho_c < q \leq v_i \leq f_\sigma(v_i)$. This implies that $\Phi(u)$ is a configuration in $[0,1]^n$ such that $\rho_c < \Phi(u)_i$ for all $i \in [n]$, and such that $q \leq \min_{i\in[n]} \Phi(u)_i$.

Now define $u^* = [q,\ldots,q]^T$ with . From previous remarks, we have that $\lim_{t\to\infty} \Phi^t(u^*) = [1\ldots 1]^T$ and that for every $t \geq 0$ and $i \in [n]$, $\Phi^t(u^*)_i \leq \Phi^t(u)_i \leq 1$. We deduce that $\lim_{t\to\infty} \Phi^t(u) = [1\ldots 1]^T$.

The case when $\rho_c > u_i$ for all $i \in [n]$ is analogous. □

In next lemma we bound the difference between the dynamic with amplification respect to the dynamic without amplification. Before giving the lemma, consider the following two remarks. First observe that, if $x \in [0,1]$,

$$-\frac{\rho_c}{4} \leq b_{\rho_c}(x) \leq \frac{(1-\rho_c)}{4}.$$

Indeed, as $x(1-x)$ reaches its maximum at $x = 1/2$. Then, we have that

$$-\frac{\rho_c}{4} \leq -\rho_c x(1-x) \leq b_{\rho_c}(x) \leq (1-\rho_c)x(1-x) \leq \frac{1-\rho_c}{4}.$$

Second, notice that since all the coordinates of matrix C are positive, then Φ_0 is monotone. More precisely, for every pair of configurations $w^1, w^2 \in [0,1]^n$, such that for all $i \in [n]$ $w_i^1 \leq w_i^2$, we have that $(Cw^1)_i \leq (Cw^2)_i$.

Lemma 2.4 *Let $u \in [0,1]^n$ be an initial configuration and $\sigma > 0$. Then for every $t \geq 0$ and $i \in [n]$,*

$$\min_{j \in [n]} (C^t u)_j - \frac{\sigma t \rho_c}{4} \leq \Phi_\sigma^t(u)_i \leq \max_{j \in [n]} (C^t u)_j - \frac{\sigma t(1-\rho_c)}{4}$$

Proof. Let $u \in [0,1]^n$ and $\sigma > 0$. We now show by induction on $t \geq 1$ that

$$(C^t u)_i - \frac{\sigma t \rho_c}{4} \leq \Phi_\sigma^t(u)$$

In the base case $t = 1$ we have that:

$$\begin{aligned} \Phi_\sigma(u)_i &= (Cu)_i + \sigma b_{\rho_c}((Cu)_i) \\ &\geq Cu_i - \frac{\rho_c \sigma}{4} \end{aligned}$$

Now, for the inductive step, suppose that the property holds for $t \geq 1$. Then,

$$\begin{aligned} \Phi_\sigma^{t+1}(u)_i &= (C\Phi_\sigma^t(u))_i + \sigma b_{\rho_c}((C\Phi_\sigma^t(u)_i) \\ &\geq (C\Phi_\sigma^t(u))_i - \frac{\rho_c \sigma}{4} \end{aligned}$$

By the induction hypothesis, we know that $\Phi_\sigma^t(u) \geq C^t u - \frac{\rho_c t \sigma}{4}[1, \ldots, 1]^T$. Then, the second remark above implies that:

$$\begin{aligned} \Phi_\sigma^{t+1}(u)_i &\geq (C(C^t u - \frac{\rho_c t \sigma}{4}[1, \ldots, 1]^T))_i - \frac{\rho_c \sigma}{4} \\ &\geq (C^{t+1} u - \frac{\rho_c t \sigma}{4} C[1, \ldots, 1]^T))_i - \frac{\rho_c \sigma}{4} \\ &\geq (C^{t+1} u)_i - \frac{\rho_c t \sigma}{4} - \frac{\rho_c \sigma}{4} \\ &\geq (C^{t+1} u)_i - \frac{\rho_c (t+1) \sigma}{4} \end{aligned}$$

We deduce that

$$\Phi_\sigma^t(u) \geq (C^t u)_i - \frac{\sigma t \rho_c}{4} \geq \min_{j\in[n]}(Cu)_j - \frac{\sigma t \rho_c}{4}$$

By analogous arguments we conclude that

$$\Phi_\sigma^t(u)_i \leq \max_{j\ \in[n]}(C^t u)_j - \frac{\sigma t(1-\rho_c)}{4}.$$

□

We are now ready to give the proof of Theorem 2.1.

Proof. (of Theorem 2.1) Without loss of generality, let us assume that $\rho_c \geq 1/2$.Since in finite dimension all norms are equivalent, we have from Lemma 2.2 that when $t \to \infty$:

$$\max_{u\in[0,1]^n} \left\| \left(C^t - \frac{U}{n} \right) u \right\|_\infty \to 0$$

this implies that there exists a $t_0 = t_0(n, \varepsilon)$ such that for every $t \geq t_0$,

$$\max_{u\in[0,1]^n} \left\| \left(C^t - \frac{U}{n} \right) u \right\|_\infty \leq \frac{\varepsilon}{3}.$$

Then, from Lemma 2.4 applied to t_0 we have that for all $i \in [n]$

$$\min_{j\ \in[n]}(C^{t_0}u)_j - \frac{\sigma t_0 \rho_c}{4} \leq \Phi_\sigma^{t_0}(u)_i \leq \max_{j\ \in[n]}(C^{t_0}u)_j - \frac{\sigma t_0(1-\rho_c)}{4}.$$

Hence,

$$\rho(u) - \frac{\varepsilon}{3} - \frac{\sigma t_0 \rho_c}{4} \leq \Phi_\sigma^{t_0}(u)_i \leq \rho(u) + \frac{\varepsilon}{3} - \frac{\sigma t_0(1-\rho_c)}{4}.$$

Thus, for every $i \in [n]$,

$$\left|\rho(u) - \Phi_\sigma^{t_0}(u)_i\right| \leq \frac{\varepsilon}{3} + \frac{\sigma t_0 \rho_c}{4}.$$

Now let us pick $\sigma = \dfrac{4\varepsilon}{3t_0\rho_c}$ in order to obtain that

$$\left|\rho(u) - \Phi_\sigma^{t_0}(u)_i\right| \leq \frac{2\varepsilon}{3}.$$

If u is such that $\rho(u) > \rho_c + \varepsilon$, then $\Phi^{t_0}(u)_i > \rho_c$ for every $i \in [n]$. In other words, after t_0 iterations of the dynamic, all the coordinates of the reached configuration has all its coordinates greater than ρ_c. From Lemma 2.3 we deduce that

$$\lim_{t\to\infty} \Phi_\sigma^t(u) = [1, \ldots, 1]^T.$$

By analogous arguments, when $\rho(u) < \rho_c - \varepsilon$ we deduce that

$$\lim_{t\to\infty} \Phi_\sigma^t(u) = [0, \ldots, 0]^T.$$

□

2.4 Experimental Results

In this section we report the results of an empirical study of the large diffusion and small amplification dynamics. Our goal is to describe how the topology of a given graph G influences the convergence time and the effectiveness of the classification for fixed values of σ (we are considering the situation where the critical density $\rho_c = 1$). We evaluate the dynamics on topologies having different densities (total number of edges), and different degree distributions.

Our study is carried out considering the following graphs classes of connected graphs.

- **Complete graphs** K_n, which correspond to the n-node graphs with all possible edges.
- **Regular graphs**, which are the n-node graphs where each node has the same degree d. In our experiments we consider $d \in \{4, 6, 8\}$.
- **Star graphs** S_n, consisting on the n-node graphs where one node has degree $n-1$ and the other $n-1$ nodes have degree 1.
- **Erdös-Rény graphs** [5], which is a model of random graphs, where each edge is independently included in the graph with probability p. For our experiments, for an n-node graph, we pick $p = \frac{2\ln(n)}{n}$, which roughly corresponds to the minimum probability that ensures that the graph is connected with high probability [5].
- **Barabási-Albert graphs**[2], which is another model of random graphs, generated sequentially according to a parameter m as follows. The network starts with m nodes connected randomly. Then, $n-m$ nodes are sequentially added to the graph. Each time that a new node is added it is connected to m existing nodes with a probability that is proportional to the degree that existing nodes already have. More formally, for each $i \in \{m+1, \ldots, n\}$, let us call $|E(G_i)|$ the total number of edges in the graph induced by $G_i = G[\{1, \ldots, i-1\}]$ and $d_i(v)$ the degree of node $v \in \{1, \ldots, i-1\}$ on G_i. Then, when node i is included in the graph it is connected to m random neighbors, where node $v \in G_i$ is picked with probability $\frac{d_i(v)}{|E(G_i)|}$. Observe that Barabási-Albert graphs with parameter m are trees.

The choice of the graph classes is based on the following criteria. First, complete graphs are used as a benchmark of ideally connected topologies, for which the local dynamics behaves in the most efficient way.

Second, the dynamics over regular graphs corresponds to the most natural generalization of the dynamics given in [7], which is designed specifically for d-dimensional tori, which are a particular case of $2d$-regular graphs.

Third, star graphs correspond to the extreme opposite of regular graphs, where one node has maximum degree, and all the others have minimum degree.

Finally, Erdös-Rényi and Barabási-Albert graphs are the best-known models of random graphs, which are commonly used to represent real-world graphs such as social networks, biological systems or particle systems.

For our simulations we pick odd values of n, from 11 to 251 with a step of 30. For each n, we pick the n-node complete graph K_n, the n-node star S_n, and, also, 50 independent random samples of:

- d-Regular graphs for $d \in \{4, 6, 8\}$ (50 samples for each d);
- Erdös-Rényi graphs with parameter $p = \frac{(2\ln(n))}{n}$;
- Barabási-Albert graphs of parameter $m \in \{1, 2, 3, 4\}$ (50 samples for each m).

If a sampled graph is not connected, it is discarded and not considered in the statistics. For each n and for each graph, $10 \cdot n$ random initial configurations are picked uniformly at random.

2.4.1 Experimental study when $\sigma = 0$

We first study the convergence time of the dynamics without amplification (i.e. when $\sigma = 0$). The dynamics is iterated until the difference between each pair of coordinates of two consecutive configurations is smaller than 10^{-3}. Finally, we take the mean convergence time over all configurations and all graphs sampled in the corresponding class.

Our empiric study of the dynamic without amplification indicates that the convergence time is affected by the edge density of the input graph, and that it is fundamentally different on trees. The results are reported in Figure 2.1.

Our results show that for complete graphs, Erdös-Rényi graphs and regular graphs, the convergence time exhibits a logarithmic growth with respect to the size of the graph. Obviously, complete graphs have the lowest mean convergence time. In the case of regular graphs, the convergence time exhibits a growth inversely proportional to the degree.

Interestingly, Erdös-Rényi graphs have the greatest mean convergence time of this group, despite the fact that that each node has an expected degree $2\ln(n)$. Also, for these graphs, the convergence time exhibits a greater standard deviation, implying that the curve is less smooth than the other ones. Nevertheless, the mean convergence time of Erdös-Rényi graphs still follows a logarithmic growth.

A second group is formed by the Barabási-Albert graphs with parameter m different than 1. For this group, the convergence time exhibits a linear growth, and it is inversely proportional to m.

Finally, a third group is formed by the star graphs and the Barabási-Albert graphs with parameter $m = 1$. Observe that these two classes are subclasses of trees. These classes exhibit a mean convergence time that is much larger than the others.

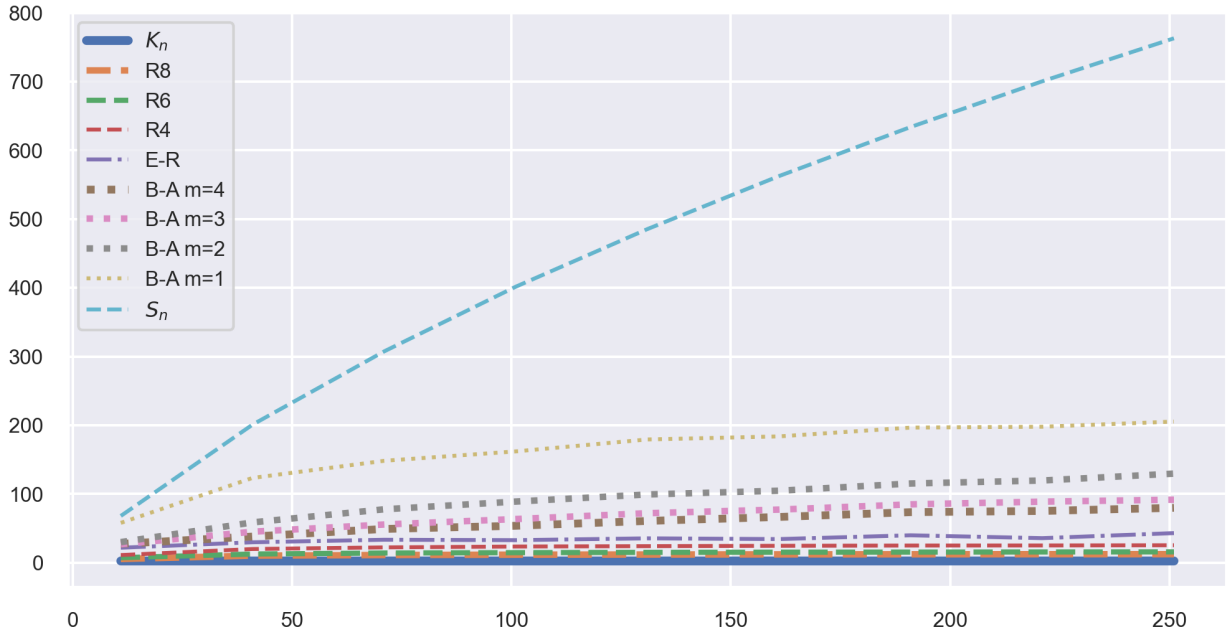

Figure 2.1: Mean convergence time of the dynamics without amplification ($\sigma = 0$). Each line represents the mean convergence time of the dynamic in the corresponding set of samples. Lines labeled K_n and S_n represent the complete and star graphs. Lines labeled R4, R6, R8 denote mean convergence time of samples of d-regular graphs for d equals 4, 6 and 8, respectively. The line labeled E-R represent the convergence time of the samples of the Erdös-Rényi graphs. Finally lines $B - Am = k$ represent the convergence time of the samples of the Barabási-Albert of parameters $k \in \{1, 2, 3, 4\}$.

2.4.2 Experimental study when $\sigma = 0.01$ and $\rho_c = 1/2$

In this subsection we report the experimental results when a small amplification is introduced. For simplicity, our study is carried out for a critical density $\rho_c = 1/2$. We empirically tested different values of σ in order to obtain a 100% of correct classification on every configuration and graph we tested. A report over the effectiveness versus correctness can be found in next subsection.

We fix $\sigma = 0.01$ and apply the same same simulation framework that we applied in the previous section. However, in this case, the convergence time is approximated by the time-step on which either every node has a state greater than 0.99 (i.e., when an all 1's configuration is approximately reached), or every node has a state smaller than 0.01 (i.e., when an all 0's configuration is approximately reached). In the first case we assume that the dynamics has classified the initial configuration with a density greater than $1/2$, while in the second case we assume that the dynamics has classified the initial configuration with a density smaller than $1/2$.

From our experimental results we deduce the following observations.

First, we obtain a dramatic difference on the convergence time for the classes that are not trees, namely complete, regular, Erdös-Rényi and Barabási-Albert graphs with parameter m

different than 1 For all these graph classes the convergence time is almost indistinguishably. In fact, the differences are roughly the same that we observed in the convergence time of the dynamics without amplification. For clarity, we show in Figure 2.2 only the convergence times for graph sizes in the rank between 221 and 251.

Second, for the graph classes that are trees, namely the star graph and the Barabási-Albert graphs with parameter $m = 1$, we observed that the convergence time diverges from the one observed for the other graph classes. This difference is especially dramatic for the Barabási-Albert graphs with parameter $m = 1$ (see Figure 2.3).

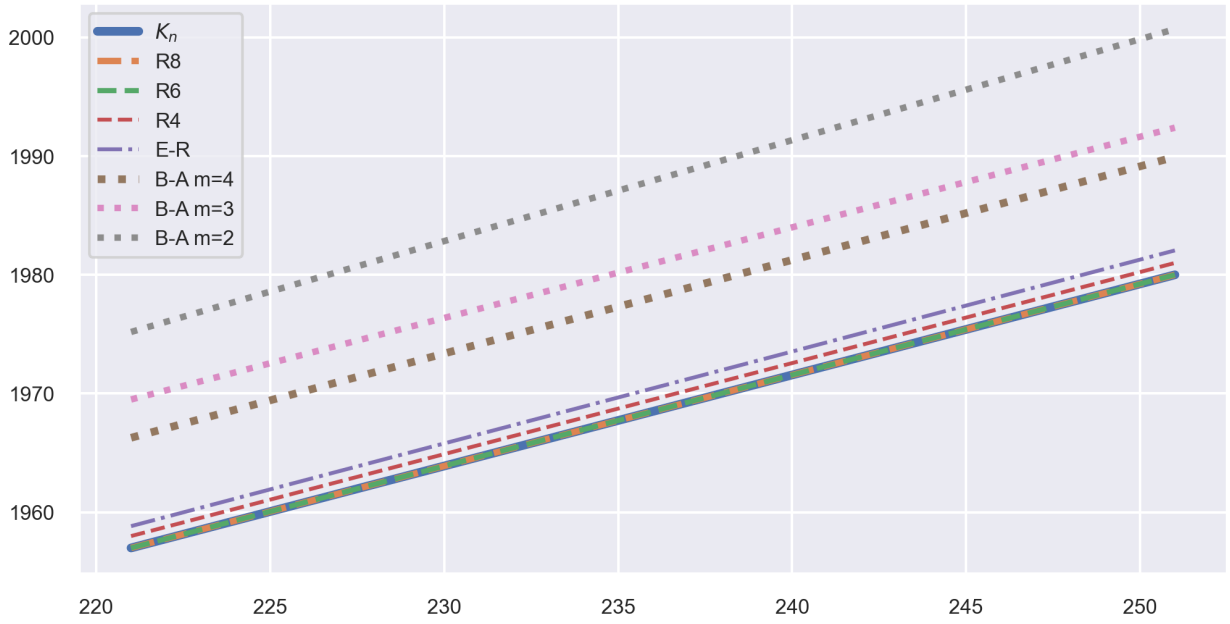

Figure 2.2: Mean convergence time of the dynamics with amplification $\sigma = 0.01$ for graph sizes between 221 and 251. Each line represents the mean convergence time of the dynamics in the corresponding set of samples. Line labeled K_n represents the complete graph. Lines labeled R4, R6, R8 denote mean convergence time of samples of d-regular graphs for d equals 4, 6 and 8, respectively. The line labeled E-R represents the convergence time of the samples of the Erdös-Rényi graphs. Finally lines B-A $m = k$ represent the convergence time of the samples of the Barabási-Albert with parameters $k \in \{2, 3, 4\}$.

2.4.3 Convergence time and effectiveness versus amplification

In this last subsection we study the influence of the amplification factor on the effectiveness and the convergence time of the dynamics. For $n = 11$ and values of σ from 0.01 to 1, we compute the number of iterations on which the dynamic reaches an (approximately) all 1's configuration or an (approximately) all 0's configuration, with the same criteria of last section. The results are reported in Figures 2.4 and 2.5.

The results for complete, regular, Erdös-Renyi, and Barabási-Albert graphs with parameter m different than 1 are very similar. In all these classes the behavior is roughly the same.

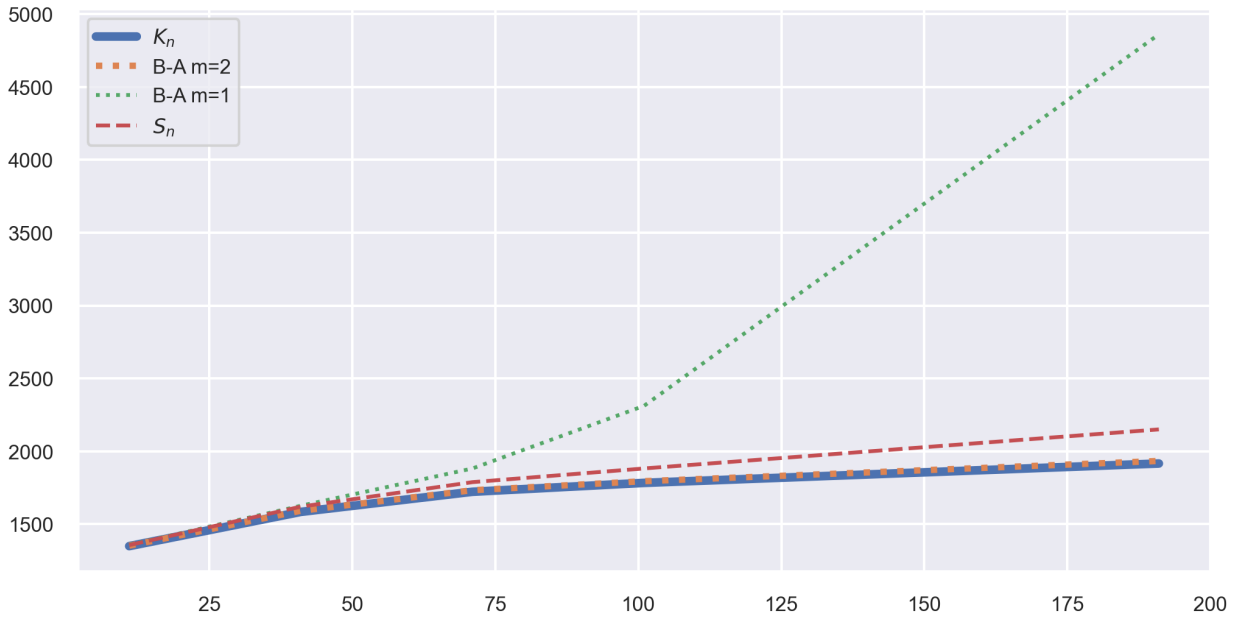

Figure 2.3: Mean convergence time of the dynamics with amplification $\sigma = 0.01$ for different graph sizes. Each line represents the mean convergence time of the dynamics in the corresponding set of samples. Lines labeled K_n and S_n represent the complete and the star graphs. Lines B-A $m = k$ represent the convergence time of the samples of the Barabási-Albert of parameters $k \in \{1, 2\}$.

For that reason, in Figure 2.4 we choose the complete graph as a representative of this set of graph classes. For $\sigma > 0.1$, the convergence time is roughly small (bounded by 200 time-steps), while for $\sigma < 0.1$ the convergence time grows exponentially. For the complete, regular and Barabási-Albert graphs with parameter $m \in \{3, 4\}$, the classification effectiveness was a 100% on every $\sigma \leq 1$. On the contrary, for the Erdös-Renyi graphs and the Barabási-Albert graphs with parameter $m = 2$, the classification effectiveness drops from 100% for $\sigma > 0.35$ (see Figure 2.5).

As we may expect from previous simulations, the behavior differs when we analyze the star graphs or the Barabási-Albert graphs with parameter $m = 1$. As for the other classes, the convergence time diverges when σ is smaller than 0.1. However, in the case of the star graph we found that the effectiveness of classification drops from 100% when $\sigma > 0.6$. Around that threshold (for $\sigma > 0.56$), the convergence time tends to augment, but it does not explodes as for small values of σ. Interestingly, the effectiveness star graphs oscillates between an 88% and a 100% as σ approaches 1 Finally, with respect to Barabási-Albert graphs with parameter $m = 1$, we observe a behavior that is similar to the star graphs, with two relevant differences. First, the classification effectiveness drops from 100% when $\sigma > 0.18$, which is smaller than the observed threshold for star graphs. Second, as σ approaches 1 the drop in the effectiveness is monotonic.

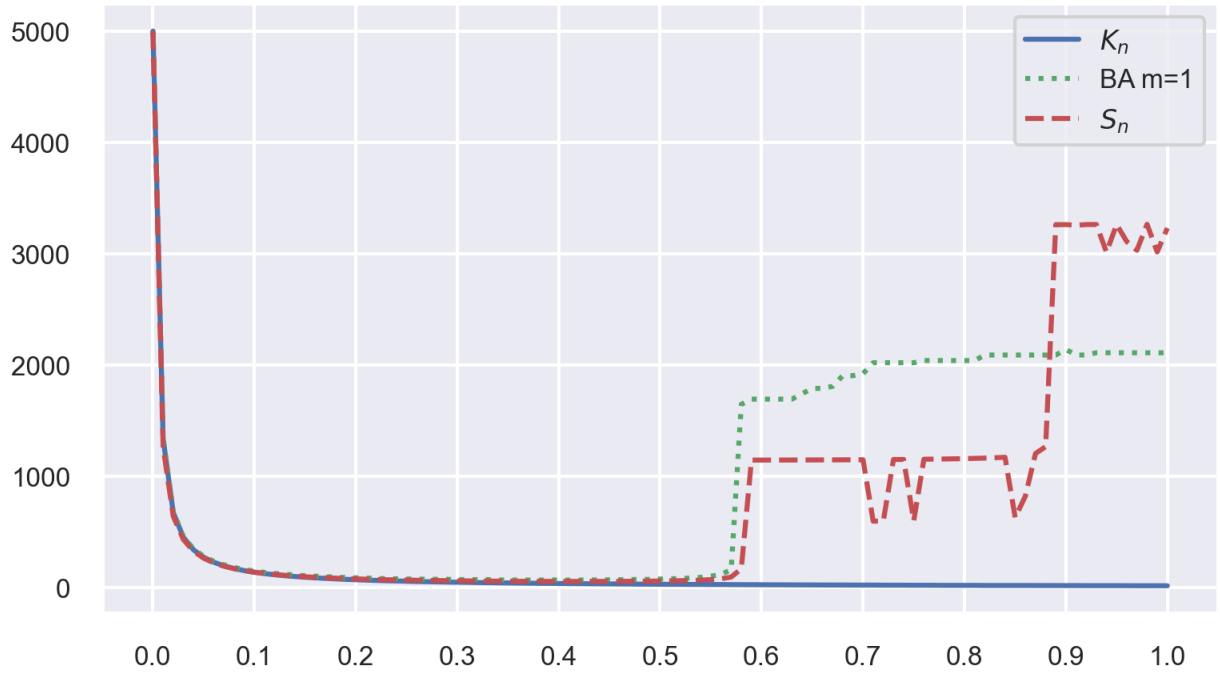

Figure 2.4: Mean convergence time for different values of σ for the 11 node complete graphs, star graphs and Barabási-Albert graphs with parameter $m = 1$.

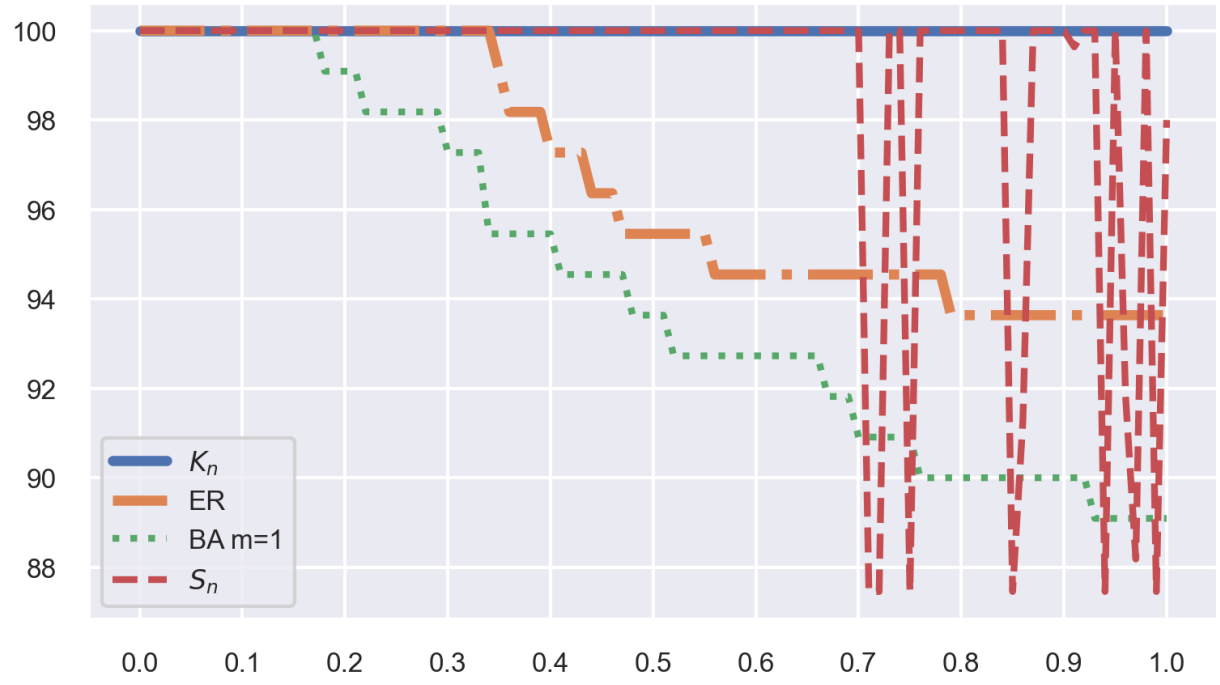

Figure 2.5: Mean percentage of effectiveness for different values of σ for the 11 node complete, Erdös-Renyi, star and Barabási-Albert graphs with parameter $m = 1$.

Chapter 3

Maximal Independent Set Computation Driven by Finite-State Dynamics

A Maximal Independent Set (MIS) is an inclusion maximal set of pairwise non-adjacent vertices. The computation of an MIS is one of the core problems in distributed computing. In this thesis, we introduce and analyze a finite-state distributed randomized algorithm for computing a MIS on arbitrary undirected graphs. Our algorithm is self-stabilizing (reaches a correct output on any initial configuration) and can be implemented on systems with very scarce conditions. We analyze the convergence time of the proposed algorithm, showing that in many cases the algorithm converges in logarithmic time with high probability.

3.1 Introduction

The Maximal Independent Set problem (MIS) is one of the main problems in distributed computing. In its simplest version, it consists in finding an inclusion-maximal set of pairwise non-adjacent vertices on an undirected graph. The MIS corresponds to a specific case of a wide problem in distributed graph algorithms, known as symmetry breaking. When a distributed algorithm is executed, the nodes of a distributed system are assumed to be in the same state, but in the successive time-steps the nodes are expected to play different roles, hence breaking the symmetry.

The MIS has a trivial solution in the classical sequential setting, where a greedy algorithm sequentially picks an arbitrary vertex, includes it in the maximal independent set, and removes that vertex together with all its neighbors. In the 80's, Karp and Wigderson [27] mentioned that the MIS is an interesting problem in non-centralized computation. Soon after that, Luby [33] and Alon, Babai, and Itai [3] presented simple distributed randomized algorithms solving MIS in $\mathcal{O}(\log n)$ time. Since then, this problem has been studied extensively in the distributed setting. In the LOCAL model, the fastest deterministic MIS algorithms for general graphs run in $\mathcal{O}(\log^5 n)$ [22], and $\mathcal{O}(\Delta + \log^* n)$ time [4]. Ghaffari [21] also obtained a

$\mathcal{O}(\log \Delta)+2^{\mathcal{O}(\sqrt{\log \log n})}$ time randomized algorithm on general graphs, and a $\mathcal{O}(\log a+\sqrt{\log n})$ time randomized algorithm for graphs of arboricity a. With respect to lower bounds, Linial [32] proved that computing an MIS on an n-cycle requires time $\Omega(\log^* n)$. Moreover, Kuhn, Moscibroda and Wattenhofer [29] showed a $\Omega(\sqrt{\log n})$ lower-bound on the round complexity on general graphs.

Another branch of research regarding the MIS problem in the distributed setting consists in considering models with limited resources. One example is the *beeping model* [1], where the nodes are limited to an extremely harsh system of communication. On each round, a node can either broadcast a signal (a beep) or hear whether a neighbor emitted a beeping signal, but noes it is not capable of distinguishing the number nor the sources of the beeping signals it receives. In this model, Afek et al. [1] showed that an MIS can be computed in time $\mathcal{O}(\log^3 n)$ when the nodes have to know the size of the graph n and have $\text{poly}(\log n)$ sized memory. The *stone-age model* is another relevant model with limited resources, where the memory of each node is limited to a constant not depending on the size of the graph. In this model, Emek and Wattenhofer [17] give an MIS algorithm with a running time of $\mathcal{O}(\log^2 n)$. Interestingly, this algorithm requires that the nodes start in a particular initial state in order to be capable of performing correct computation.

An algorithm is called *self-stabilizing* [25] if it can reach a correct output starting from any initial state. The motivation for this kind of algorithm is the capacity of distributed systems to self-repair a faulty configuration when one of the parties crashes. For instance, consider a maximal independent set, where one of the nodes crashes. It is possible that the rest of the nodes do not form a maximal independent set in the remaining graph (for instance, consider an MIS in a complete graph where the unique marked node crashes). In that context, a self-stabilizing MIS algorithm should be able to reach an MIS for any initial state configuration of the nodes.

In a keynote talk of SIROCCO 2022 [23], George Giakkoupis presented two extremely simple randomized algorithms for MIS. His algorithms have two interesting properties: they are self-stabilizing and they require only two or three states. Despite this simplicity, the algorithm has hardly been studied before, and its convergence time is not been settled yet. In this article, we propose a variant of the algorithm of [23] and study its convergence time both numerically and analytically.

Notation. For a positive integer k, we denote by $[k]$ the set $\{1, \dots, k\}$. Also, for a set S, we denote by $x \in_U S$ the process of taking an element of S uniformly at random. On inputs of size n, we say that an event occurs with high probability if it occurs with probability greater or equal than $1 - 1/n$.

The dynamics. Let $G = (V, E)$ be a simple finite undirected graph and $k \geq 2$ an integer. A *configuration* is a function that assigns to each node a *state* in $[k]$. Formally, a configuration is given by a function $x : V \to \{0\} \cup [k]$. For each node $u \in V$ we denote x_u the state of u on configuration x. The nodes in a state different than zero are called *marked nodes*, while the nodes in state 0 are called *unmarked*.

We say that a configuration *represents an independent set of* G when no edge has both

endpoints in a state different than 0. Formally,

$$\forall \{u, v\} \in E,\ x_u \cdot x_v = 0.$$

Additionally, we say that a configuration *represents a maximal independent set of* G when it represents an independent set and, in addition, no edge has both endpoints in state 0. Formally,

$$\forall u \in V,\ x_u \neq 0 \vee (\exists v \in N(u),\ x_u \neq 0)\,.$$

Let us consider the following stochastic dynamic over graph configurations. Given a configuration x, the next configuration x' is computed synchronously according to a rule that we denote MIS-Dynamics: Synchronously, the new state x'_u of $u \in V$ is computed as follows:

(i) If $x_u = 0$ and $\forall v \in N(u), x_v = 0$, then $x'_u \in_U [k]$.

(ii) If $x_u \neq 0$ and $\exists v \in N(u), x_u = x_v$ and $\forall v \in N(u), x_v \leq x_u$, then $x'_u \in_U [k]$.

(iii) If $x_u \neq 0$ and $\exists v \in N(u), x_v > x_u$, then $x'_u = 0$.

(iv) $x'_u = x_u$ otherwise.

In the special case of two-state configurations $x : V \to \{0, 1\}$, we consider the following stochastic dynamics, that we call 2-MIS-Dynamics: Synchronously, we compute for each $u \in V$ the new state x'_u as follows:

(i) If $x_u = 0$ and $\forall v \in N(u), x_v = 0$, then $x_u = 1$.

(ii) If $x_u \neq 0$ and $\exists v \in N(u), x_v \neq 0$, then $x_u \in_U \{0, 1\}$.

(iii) $x'_u = x_u$ otherwise.

Given a configuration x (also called initial configuration), the *trajectory* of x, denoted $\{x^t\}_{t \geq 0}$, is the random variable representing the evolution of the MIS-Dynamics where x^t is obtained from x^{t-1} for each $t > 1$, and $x^0 = x$.

We say that a configuration is a *fixed point* for the MIS-Dynamics if $x' = x$ with probability 1. Observe that a configuration x is a fixed point of the MIS-Dynamics if and only if x represents a maximal independent set. Indeed, let x be a fixed point of the MIS-Dynamics. We say that a node u is *stabilized* on x if one of the following conditions is satisfied:

1. $x_u \neq 0$ and every $v \in N(u)$ satisfies $x_v = 0$,
2. $x_u = 0$ and there exists a stabilized neighbor of u such that $x_u \neq 0$.

3.1.1 Our Contribution

We study the convergence time of the MIS-Dynamics and 2-MIS-Dynamics both empirically and analytically.

We first simulate the dynamics on different graph classes, namely: Complete graphs, stars, graphs of Erdös-Renyi, random trees, and graphs of bounded degeneracy. On the one hand, we observe that the 2-MIS-Dynamics converges to a MIS in logarithmic time in expectation for all classes. We also observe that the convergence time tends to increase with the density of the input graph. In fact, for complete graphs, we the convergence time is $\Omega(\log^2(n))$ with non-negligible probability. On the other hand, we observe that the MIS-Dynamics converges to an MIS in logarithmic time, both in expectation and with high probability. Moreover, the computation time decreases as the number of states augments.

Then, we analytically study the convergence time of the MIS-Dynamics. First, we show that with high probability, the dynamic converges in time $\mathcal{O}(\alpha \log n)$, where α is the size of a maximum independent set of the input graph. Then, we extend our analysis to the 2-MIS-Dynamics, showing that with high probability it converges in time $\mathcal{O}(\alpha \log^2 n)$. Finally, we show that restricted to the class of d-degenerate graphs, the $2-\mathsf{MIS}-\mathsf{Dynamics}$ converges in time $\mathcal{O}(\log n)$ with high probability.

3.1.2 Structure of the article

We begin giving some background and preliminaries in Section 3.2. In Section 3.3 we report the results of our computational simulations. Then, in Section 3.4 we give bounds for the convergence time of the MIS-Dynamics and 2-MIS-Dynamics on arbitrary graphs. In Section 3.5 we study the 2-MIS-Dynamics on graphs of bounded degeneracy. We finish with a discussion in Section 4.2.

3.2 Preliminaries

In this article, all graphs are simple, finite, and undirected. Given a node v of a graph $G = (V, E)$, the *neighborhood of* v, denoted by $N(v)$, is the set of vertices adjacent to v. Formally $N(v) = \{u \in V : \{u, v\} \in E\}$. The degree of a node u, denoted $d(u)$ is the cardinality of $N(v)$. Given a set of nodes $U \subseteq V$, the *subgraph of* G *induced by* U, denoted $G[U]$, is the graph defined by the vertex set U, and all the edges in E with both endpoints in U. A graph is called *connected* if there is a path between every pair of vertices. A *connected component* of a graph G is an inclusion maximal connected set of vertices. A connected graph without cycles is called a *tree*. A graph where every node has degree d is called d-regular.

A set of nodes S is called an *independent set* if the graph induced by it has no edges. An inclusion-maximal independent set is simply called *maximal independent set*. The cardinality of an independent set of maximal cardinality is denoted $\alpha(G)$, and is called the *independence number* of G. The problem of computing $\alpha(G)$ is NP-Hard. The range of its value goes from

1 for the complete graph, to $n-1$ for the case of the star graph. There are a number of combinatorial lower bounds for $\alpha(G)$ with respect to some graph parameters. In this article, we make use of the following simple result.

Proposition 3.1 ([10]) *Let G be an n-node graph of maximum degree Δ. Then, $\alpha \geq n/\Delta$.*

A graph property is called an *hereditary property* if it is closed under taking induced subgraphs. An example of a hereditary property is bounded degeneracy. A graph G is called *d-degenerate* if every subgraph of G (including G itself) contains a vertex of degree at most d. Alternatively, a graph has degeneracy d if it can be decomposed successively removing vertices of degree at most d. In the following lemma, we now show that in a d-degenerate graph, for most vertices their degree is bounded by $4d-2$.

Lemma 3.2 *A connected graph of degeneracy d contains at least $n/2$ nodes of degree at most $4d-2$.*

Proof. First, observe that a graph of degeneracy d has at most dn edges. Let us call U the set of nodes of a degree greater or equal than $4d-1$. Then, for the hand-shaking lemma:

$$2dn = \sum_{v \in V} d(v) \geq |U|(4d-1) + (n - |U|).$$

The previous bound implies that $|U| \leq \frac{(2d-1)n}{4d-2} \leq n/2$. We deduce that the cardinality of the set vertices of degree at most $4d-2$ is at least $n/2$. □

The graphs of Erdös-Rényi-Gilbert (in the following graphs of Erdös-Rényi for simplicity) are a randomized model of graphs where a graph is constructed by connecting labeled nodes, where each edge is included in the graph with probability p, independently from every other edge. The following result states that with high probability the independence number of a graph of Erdös-Rényi is bounded by the logarithm of the number of nodes.

Proposition 3.3 *[6] For each $p \in (0,1)$ and sufficiently large n, the independence number of a graph of Erdös-Rényi with parameters n and p is $\mathcal{O}(\log n)$ with high probability.*

3.3 Experimental Results

In this section, we report the empirical analysis of the 2-MIS-Dynamics and the MIS-Dynamics.

In Figure 3.1 we show the results of a study of the 2-MIS-Dynamics and the MIS-Dynamics of k states with $k \in \{2, 3, 4, 9\}$ over complete graphs. We observe that on all cases the average convergence time has a logarithmic grow with respect to the size of the graph. We also observe that the average convergence time of the MIS-Dynamics decreases as the number of states k grows. The 2-MIS-Dynamics has a better convergence time than the three-state MIS-Dynamics for small values of n, but for complete graphs on more than 16 nodes the

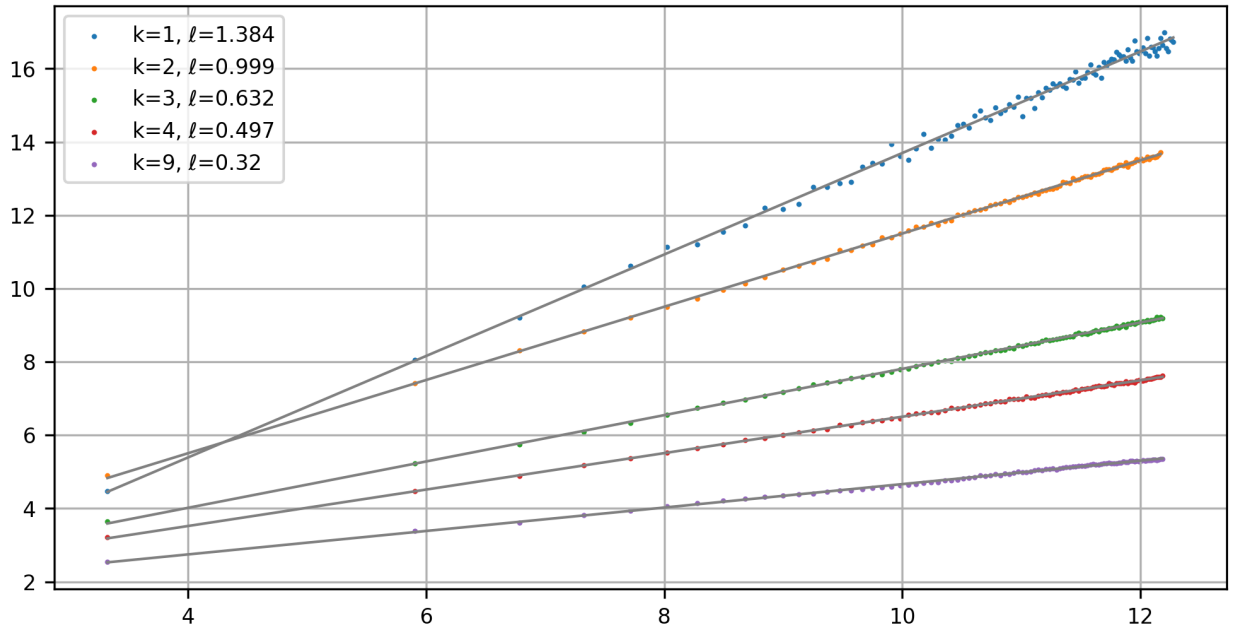

Figure 3.1: Plot of the average convergence time for the 2-MIS-Dynamics and MIS-Dynamics. On the x-axis we have the number of nodes in $\log_2$ scale, while in the y-axis are given the number of iterations. The different dynamics are represented with different colors. The line labeled $k = 1$ corresponds to the 2-MIS-Dynamics while the other lines represent the MIS-Dynamics with $k \in \{2, 3, 4, 9\}$ states. Each line represents the average convergence time of 5000 initial configurations picked uniformly at random. We also give the slope ℓ of the least squares regression line corresponding to the points.

2-MIS-Dynamics has a larger average convergence time than all the MIS-Dynamics for every $k \geq 2$.

We then extend the analysis to other graph classes. In Figure 3.2 we report our results for complete, star, random trees (Barabasi-Albert graphs with $m = 1$) and 4-regular graphs. In all cases, the results are roughly the same than those we obtained for the complete graphs. That is to say, the average convergence time $\mathcal{O}(\log n)$ and it decreases as we increase the number of states in the dynamics. We complement our analysis with a study of the worst-case convergence time (i.e. the maximum convergence time over all initial configurations). In Figure 3.3 we report our results, which were obtained in the same way as Figure 3.2, except that we take the maximum over all observed convergence times for each graph size.

Interestingly, in the worst-case analysis on complete graphs, we observe that the convergence time of the 2-MIS-Dynamics behaves significantly different than on the rest of the classes, and also with respect to the average case. In fact, our results suggest that in the worst case the convergence time of the 2-MIS-Dynamics is $\mathcal{O}(\log^2 n)$. This fact is indeed verified in the next section.

Finally, we explore how the density of the input graph influences the average and worst-case convergence-time of the 2-MIS-Dynamics and MIS-Dynamics. To do so, we fixed a number of nodes to $n = 500$, and simulate the dynamics in two families of graphs. First, on Erdös-

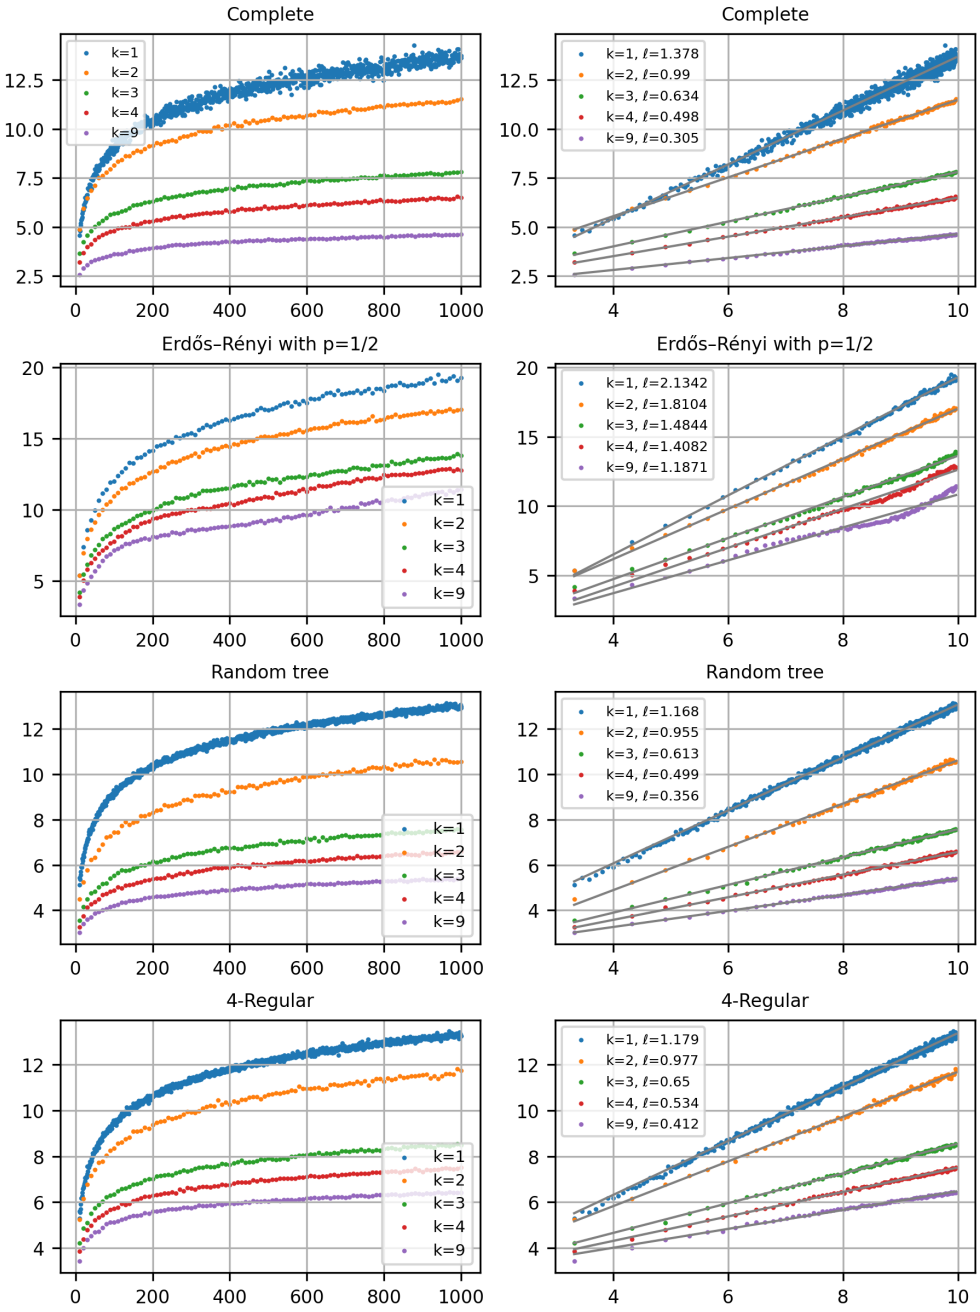

Figure 3.2: Plots of the average convergence time for the 2-MIS-Dynamics and MIS-Dynamics for different graph classes. On the x-axis we have the number of nodes, while on the y-axis are given the number of iterations. In the left column, we have the linear scale, while in the right column we give the x-axis in $\log_2$ scale. The different dynamics are represented in different colors. The line labeled $k = 1$ corresponds to the 2-MIS-Dynamics while the other lines represent the MIS-Dynamics with $k \in \{2, 3, 4, 9\}$ states. Each line represents the average convergence time of 1000 initial configurations picked uniformly at random. We also give the slope ℓ of the least squares regression line corresponding to the points.

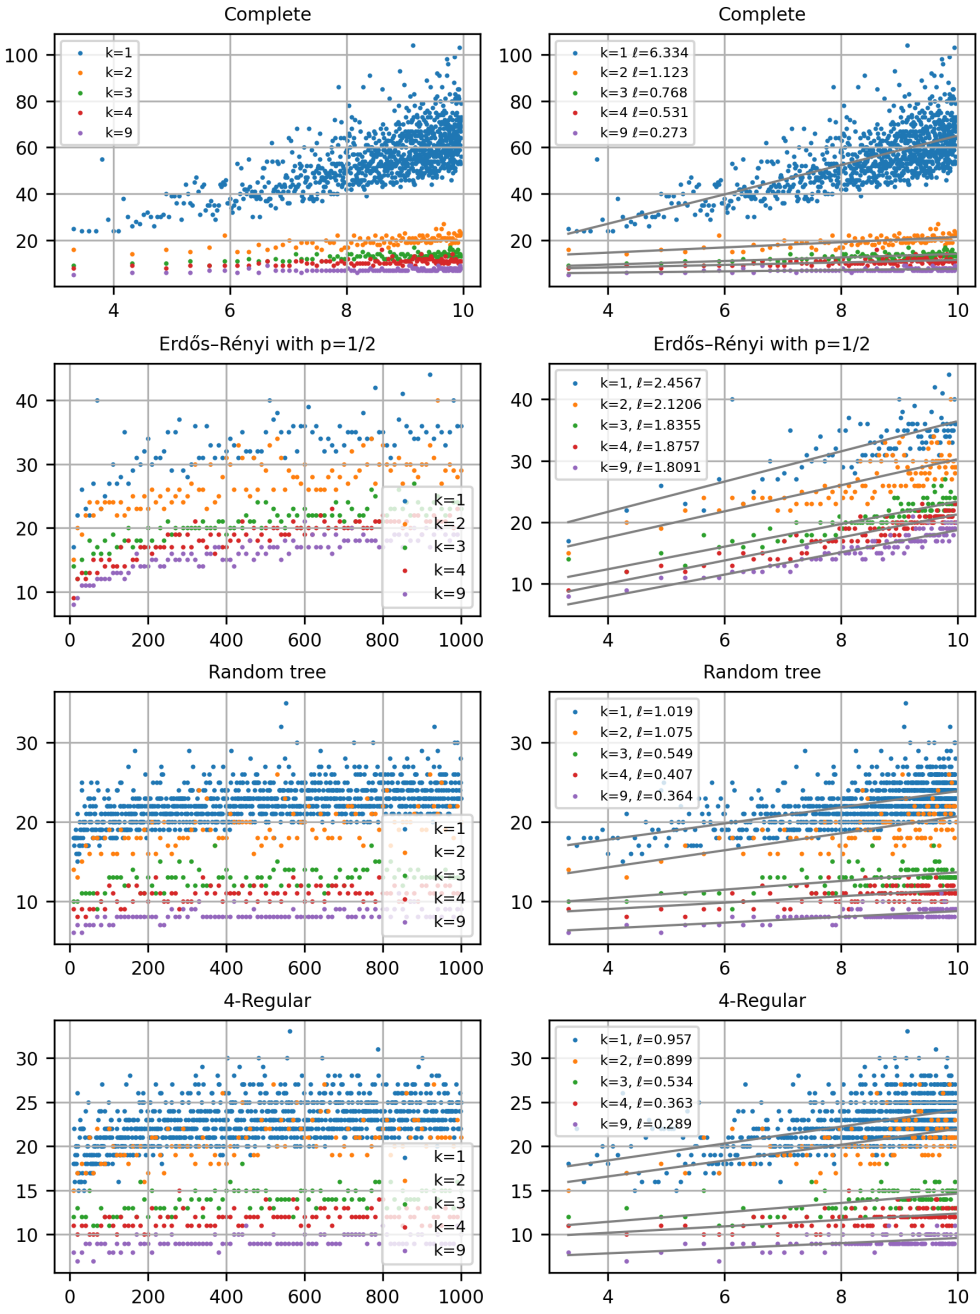

Figure 3.3: Plots of the worst-case convergence time for the 2-MIS-Dynamics and MIS-Dynamics for different graph classes. On the x-axis we have the number of nodes, while on the y-axis are given the number of iterations. In the left column, we have the linear scale, while in the right column we give the x-axis in $\log_2$ scale. The different dynamics are represented in different colors. The line labeled $k = 1$ corresponds to the 2-MIS-Dynamics while the other lines represent the MIS-Dynamics with $k \in \{2, 3, 4, 9\}$ states. Each line represents the average convergence time of 1000 initial configurations picked uniformly at random. We also give the slope ℓ of the least squares regression line corresponding to the points.

Rényi graphs for different probabilities p. Second, in d-degenerate graphs for different values of d. Our results are reported in Figure 3.4. We observe that for the 2-MIS-Dynamics and the MIS-Dynamics, the maximum average convergence times are found on graphs with a mean degree around $n/2$. With respect to the worst-case convergence-time, we observe a similar behavior for the MIS-Dynamics, while for the 2-MIS-Dynamics tends to increase with the density.

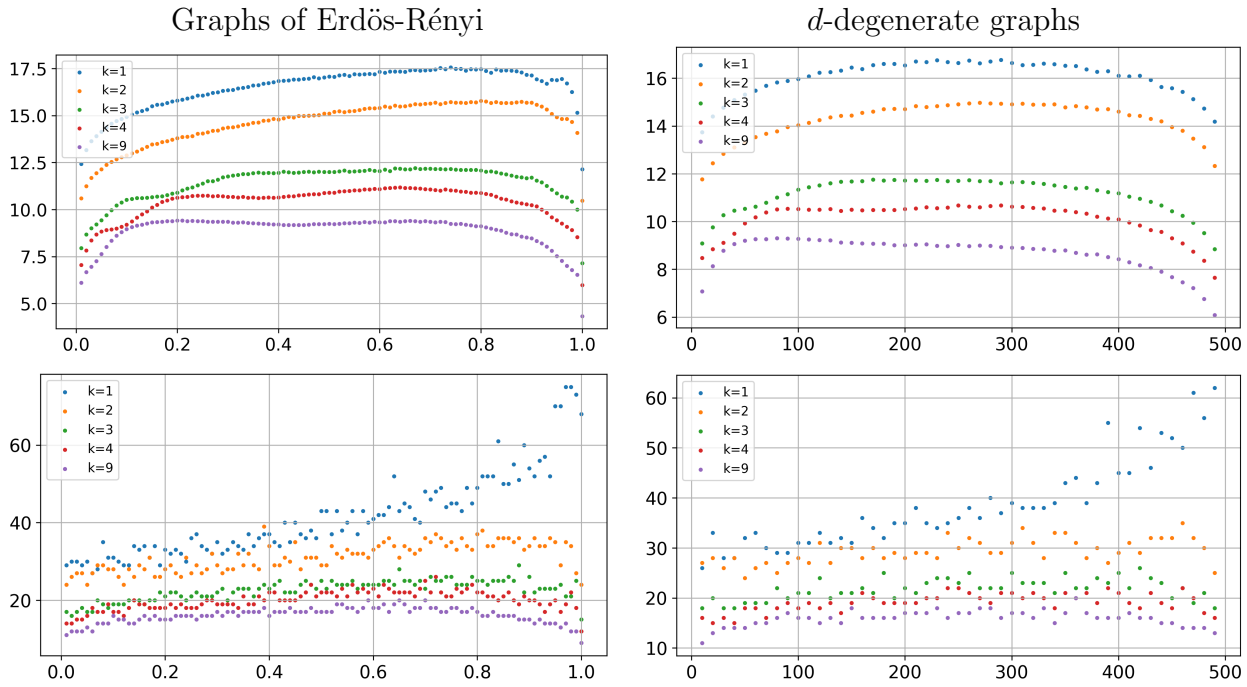

Figure 3.4: Plots of the influence of the edge density on convergence time for the 2-MIS-Dynamics and MIS-Dynamics. On the four plots we represent the convergence time of the 2-MIS-Dynamics(k=1) and the MIS-Dynamics on $k \in \{2, 3, 4, 9\}$ states on graphs with $n = 500$ nodes. In the left column we represent the behavior of the convergence-times on Erdös-Renyi graphs, where the x-axis represents the probability p of adding the edge between two vertices. In the right column we represent the behavior of the convergence-times on d-degenerate graphs. In that case, the x-axis represents the different values of the degeneracy d. In the top row the y-axis represent the average convergence times of the corresponding dynamics, while the bottom row we show the worst-case convergence times. Each point is computed taking 10000 random initial conditions, each one on a different graph picked uniformly at random.

3.4 A bound on the convergence time of the MIS-Dynamics on arbitrary graphs

In this section, we show that for every initial configuration, the MIS-Dynamics converges to a fixed point (hence a configuration that represents a maximal independent set) in $\mathcal{O}(\alpha \cdot \log n)$ time-steps on average. For simplicity, our analysis focuses on the case where $k = 2$, as it can be trivially generalized to the case where $k > 2$.

For a configuration x, we define the following *energy functional*:

$$S(x) = \sum_{\{u,v\}\in E} \delta(x_u, x_v), \text{ where } \delta(a,b) = \begin{cases} 1 & \text{if } a = b \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Let us fix an initial configuration x^0 and a time-step t. We denote by S^t the random variable that equals $S(x^t)$. A configuration x^t such that $S^t = 0$ is called a *zero-energy configuration*. Observe that all configurations that represent an independent set are zero-energy configurations (but the converse is not true). We call E^t_m the random variable representing the set of edges of G having on time-step t with both endpoints marked with the same value. Formally,

$$E^t_m = \{\{u,v\} \in E : x^t_u = x^t_v \neq 0\}$$

Observe that $S^t = |E^t_m|$. We also denote by A^t and B^t the random variables representing the following sets of edges:

$$A^t = \{u \in V : x^{t-1}_u = 0 \text{ and } x^t_u \neq 0\}$$

$$B^t = \{u \in V : x^{t-1}_u \neq 0 \text{ and } x^t_u \neq 0\}$$

We say that a time-step t such that $A^t \neq \emptyset$ is a *marking time step*.

Lemma 3.4 *Let $\{x^t\}_{t\geq 0}$ be a trajectory and let t_0 be a marking time-step. Let C be a connected component of the graph induced by A^{t_0}. Then C contains a node that is never unmarked. Formally, there is $u \in C$ satisfying $x^t_u \neq 0$ for every $t \geq t_0$.*

Proof. Let us suppose by contradiction that there is a time step $t_1 > t_0$ where every node in C has visited state 0 at least once on some configuration of $\{x^{t_0}, \ldots, x^{t_1}\}$. From all possible choices of t_1, we pick the minimum one. Let v be a node such that $x^t_v \neq 0$ for every $t_0 \leq t < t_1$ and $x^{t_1}_v = 0$. Then, the only possibility is that $x^{t_1-1}_v = 1$ and that v has a neighbor $u \in N(v)$ such that $x^{t_1-1}_u = 2$. Observe that $x^{t_1}_u \neq 0$ (it impossible that a node in state 2 switches to 0 in the next time-step). Then, by definition of T, there must exist a time-step $t_0 < t < t_1 - 1$ such that $x^t_u = 0$. However, we are assuming that v is in a state different than 0 on all the configurations of that interval. Therefore, it is impossible that u switches to a state different than 0 on a time-step between t and $t_1 - 1$. This contradicts the choice of u. We deduce that at least one node of C is never unmarked. □

Lemma 3.5 *The trajectory of every configuration has at most α marking time steps.*

Proof. Now let us call T the set of time-steps t such that $A^t \neq \emptyset$. Given a time-step $t \in T$, we know by Lemma 3.4 that there is a node $v(t) \in V$ that is never unmarked after time-step t. We have that the set $\{v(t)\}_{t\in T}$ forms an independent set of G. Indeed, let us pick two different time-steps $t_1, t_2 \in T$ such that $t_1 < t_2$. By definition of A^{t_2} we have that $x^{t_2-1}_{v(t_2)} = 0$ and $x^{t_2}_{v(t_2)} \neq 0$. Since $t_1 < t_2$ we have that $x^{t_2-1}_{v(t_1)} \neq 0$. We deduce that $v(t_1)$ and $v(t_2)$ cannot be adjacent. We conclude that $|T| \leq \alpha(G)$. □

Lemma 3.6 *Let $t > 0$ be a non-marking time-step. Then, $\mathbb{E}(S^t) \leq S^{t-1}/2$.*

Proof. First, we observe that, since $A^t = \emptyset$, every edge that contributes to S^t also contributes to S^{t-1}. Formally, $E^t_m \subseteq E^{t-1}_m$. Indeed, let $e = \{u, v\}$ be an edge contained in E^t_m. Clearly u and v belong to B^t, as A^t is empty. Then, $x^{t-1}_u \neq 0$ and $x^{t-1}_v \neq 0$. Moreover, $x^{t-1}_u = x^{t-1}_v$, as otherwise at least one of the endpoints would be 0 on time-step t. Therefore e is also contained in E^{t-1}_m. Now observe that for every $e \in E^{t-1}_m$, the probability that $e \in E^t_m$ is fewer or equal than the probability that both endpoints of e remain marked and choose the same state, which is $1/2$. We deduce that $\mathbb{E}(S^t) \leq S^{t-1}/2$. □

Lemma 3.7 *Let t be a zero energy time-step. Then at least one of the following holds:*

- *$t + 1$ is marking,*
- *$t + 2$ is marking,*
- *x^{t+1} is a fixed point.*

Proof. Let t be a time-step satisfying that $S^t = 0$, and let us assume that $A^{t+1} = \emptyset$. We show first that x^{t+1} must represent an independent set. Observe that $A^{t+1} = \emptyset$ implies $S^{t+1} = 0$. Then, in x^t or x^{t+1} no edge has both endpoints in state different than 0. Suppose that there exist an edge $\{u, v\}$ such that $x^{t+1}_u = 1$ and $x^{t+2}_v = 2$. Then necessarily $x^t_u = x^t_v = 0$, which contradicts the assumption of $A^{t+1} = \emptyset$. We deduce x^{t+1} represents an independent set. Now suppose that x^{t+1} does not represent a maximal independent set, that is to say, there is a node w such that $x^{t+1}_v = 0$ for all $v \in N(w) \cup \{w\}$. Then necessarily $x^{t+2}_w = 1$, implying that $A^{t+2} \neq \emptyset$. □

Now we are ready to prove the main result of this section.

Theorem 3.8 *For every initial configuration, the* MIS-Dynamics *converges to a configuration representing a maximal independent set in $\mathcal{O}(\alpha \cdot \log n)$ time-steps with high probability.*

Proof. Let x^0 be an arbitrary initial configuration. From Lemma 3.6 we know that, with high probability, in at most $\mathcal{O}(\log n)$ time-steps the trajectory visits a zero-energy configuration or a marking time-step. Let $t > 0$ be a marking time-step or a time-step where the trajectory visits a zero-energy configuration. Combined with Lemma 3.7 and 3.4 we know that either x^{t+1} is a fixed point (so we are done), or either t, $t+1$ or $t+2$ is marking. In the later cases, we repeat the analysis by taking x^t, x^{t+1} or x^{t+2} as the initial configuration. By Lemma 3.5 we know that the number of repetitions is bounded by α. We deduce that with high probability the dynamic converges to a configuration representing a maximal independent set in $\mathcal{O}(\alpha \cdot \log n)$ time-steps. □

3.4.1 The 2-MIS-Dynamics on arbitrary graphs

We now adapt our result to the 2-MIS-Dynamics. The difference is found in the result given on Lemma 3.4. In fact, that lemma does not hold for the 2-MIS-Dynamics. Indeed, on

every marking time-step t, there exists a non-zero probability that all nodes of A^t become unmarked. Nevertheless, the next lemma shows that at least one marked node is stabilized with a constant probability.

Lemma 3.9 *Let $\{x^t\}_{t\geq 0}$ be a trajectory and let t_0 be a marking time-step. Let C be a connected component of the graph induced by A^{t_0}. Let $\mathcal{E}$ the event where at least one node of C is never unmarked. Formally,*

$$\mathcal{E} : \exists u \in C, \forall t \geq 0, x_u^t \neq 0.$$

Then, there exists an absolute constant $c \in (0,1)$ satisfying that $Pr(\mathcal{E}) \geq c$.

Proof. Let m be the number of edges in $G[C]$. If $m = 0$ then C consists on an isolated marked node, which is stabilized with probability 1. If $0 < m \leq 144$ we have that $Pr(\mathcal{E}) \geq 2^{-144}$. Indeed, when $m \leq 144$ we have that C contains at most 144 nodes. A lower-bound on $Pr(\mathcal{E})$ is the event where on one time-step all except one fixed node of C becomes 0. The probability of such an event is lower-bounded by 2^{-144}. In the following, we assume that $m > 144$.

Let us denote $e_1, \ldots, e_m$ the edges of $G[C]$. For each $i \in [m]$ we denote by e_i^t the random variable that equals 1 if edge e_i is marked on all time-steps in $\{t_0, \ldots, t_0 + t\}$. We also denote $m^t = \sum_{i\in[m]} e_i^t$. In words, m^t is the random variable representing the edges that have both endpoints marked on all time-steps in $\{t_0, \ldots, t_0 + t\}$. Observe that for each $i, j \in [m]$, $\mathbb{E}(e_i^t) = 2^{-2t}$ and by linearity of the expectation, $\mathbb{E}(m^t) = 2^{-2t}m$.

We aim to bound the probability that the actual value of m^t has a large gap with respect to its expectation. By Chebyshev's inequality we have that

$$Pr(|m^t - \mathbb{E}(m^t)| > a) \leq \frac{\mathrm{Var}(m^t)}{a^2}.$$

By Bienyamé's identity, we know that

$$\mathrm{Var}(m^t) = \sum_{i\in[m]} \mathrm{Var}(e_i^t) + \sum_{i\in[m]} \sum_{j\in[m]\setminus\{i\}} \mathrm{Cov}(e_i^t, e_j^t)$$

where

$$\mathrm{Var}(e_i^t) = \mathbb{E}((e_i^t)^2) - (\mathbb{E}(e_i^t))^2 = 2^{-2t} - 2^{-4t}.$$

and

$$\mathrm{Cov}(e_i^t, e_j^t) = \mathbb{E}(e_i^t e_j^t) - \mathbb{E}(e_i^t)\mathbb{E}(e_i^t) = \begin{cases} 2^{-3t} - 2^{-4t} & \text{if } e_i \cap e_j \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

Then, if we denote by Δ the maximum degree of $G[C]$,

$$\mathrm{Var}(m^t) \leq m \cdot (2^{-2t} - 2^{-4t}) + m \cdot 2\Delta \cdot (2^{-3t} - 2^{-4t}).$$

Observe that $\Delta^2 \leq |C|\Delta \leq 2m$. Therefore,

$$\mathrm{Var}(m^t) \leq m \cdot (2^{-2t} - 2^{-4t}) + 2^{3/2}m^{3/2} \cdot (2^{-3t} - 2^{-4t}).$$

Now let us pick $k \leq m$ a variable to be fixed later. If we choose

$$\tau = \left\lfloor \frac{\log(m) - \log(k)}{2} \right\rfloor,$$

we have that $\mathbb{E}(m^\tau) \in [k, 2k]$ and

$$\begin{aligned}\mathrm{Var}(m^\tau) &\leq m \cdot \left(\frac{k}{m} - \left(\frac{k}{m}\right)^2\right) + 2^{3/2} m^{3/2} \cdot \left(\left(\frac{k}{m}\right)^{3/2} - \left(\frac{k}{m}\right)^2\right) \\ &= k + k\sqrt{8k} - \left(\frac{k^2}{m} - \frac{\sqrt{8}k^2}{\sqrt{m}}\right) \leq k + k\sqrt{8k}\end{aligned}$$

Then,

$$Pr(|m^\tau - \mathbb{E}(m^\tau)| > k/2) \leq \frac{4}{k} + \frac{4\sqrt{8}}{\sqrt{k}}.$$

If we pick $k = 144$ we obtain that $Pr(|m^\tau - \mathbb{E}(m^\tau)| > 72) < 0.98$. In other words, with probability greater than 0.02 we have that $m^\tau \in [\mathbb{E}(m^\tau) - 72, \mathbb{E}(m^\tau) + 72] \subseteq [72, 360]$. Conditioned to that event, we have that with probability greater or equal than 2^{-360} a fixed node u is stabilized on time-step $\tau + 1$.

We conclude that $Pr(\mathcal{E}) \geq c = 0.02 \cdot 2^{-360}$. □

Using the previous lemma, we can show that, with high probability, there are $\mathcal{O}(\alpha \cdot \log n)$ marking time-steps on a 2-MIS-Dynamics.

Lemma 3.10 *For every initial configuration, with high probability, there are $\mathcal{O}(\alpha \cdot \log n)$ marking time-steps on the* 2-MIS-Dynamics*.*

Proof. From Lemma 3.9, we know that there exists a constant c such that at least one marked node is stabilized with a probability greater than c. Then, on $\mathcal{O}(\log n)$ marking time-steps at least one node is stabilized with high probability. We deduce that the trajectory of every initial configuration visits $\mathcal{O}(\alpha \cdot \log n)$ marking time-steps with high probability. □

Observe that we can show results analogous to Lemmas 3.6 and 3.7 using exactly the same proofs. We deduce the main result of this subsection.

Theorem 3.11 *For every initial configuration, the* 2-MIS-Dynamics *converges to a configuration representing a maximal independent set in $\mathcal{O}(\alpha \cdot \log^2 n)$ time-steps with high probability.*

3.5 The 2-MIS-Dynamics on graphs of bounded degeneracy

In this section, we show that the 2-MIS-Dynamics converges to a configuration that represents a maximal independent set in time $\mathcal{O}(\log n)$ on average.

Let G be an arbitrary graph. We denote by $G_{\leq d}$ the sub-graph of G induced by the nodes of degree at most d and by α_d the size of a maximum independent set of $G_{\leq d}$.

Lemma 3.12 *Let G be a graph, let d be a positive integer, and let x be a configuration of G with no stabilized vertices. Then, on average, $\Omega(\frac{\alpha_d}{4^{d+1}})$ nodes are stabilized after two time-steps.*

Proof. For each $u \in V$ and $t > 0$, let us call $P(u,t)$ the probability that u is stabilized on time-step t. We claim that $P(u,2) \geq 2^{-(2d_u+1)}$. Indeed, if $x_u = 1$ then $P(u,1) \geq 2^{-(d_u+1)}$, hence $P(u,2) \geq 2^{-(2d_u+1)}$. If $x_u = 0$ then, on time step $t = 1$ the probability that u and all its neighbors are unmarked is at least 2^{-d_u}. Hence $P(u,2) \geq 2^{-(2d_u+1)}$.

Now let U be a maximum independent set of $G_{\leq d}$, and let W be the random variable representing the subset of U that is stabilized on time-step $t = 2$. Then, for each $u \in U$, $Pr(u \in W) \geq 2^{-(2d+1)}$. Hence

$$\mathbb{E}(|W|) \geq \frac{|U|}{2^{2d+1}} \geq \frac{\alpha_d}{4^{d+1}}.$$

□

Theorem 3.13 *For every initial configuration over a d-degenerate graph, the* 2-MIS-Dynamics *converges to a configuration representing a maximal independent set in $\mathcal{O}(\log n)$ time-steps with high probability.*

Proof. Let G be an arbitrary n-node graph of degeneracy d, and let x be an arbitrary configuration of G. Without loss of generality, we assume that G has no stabilized vertices. Otherwise, we pick the set of nodes U that are not stable on x, and continue the reasoning with the subgraph of G induced by U. Observe that $G[U]$ is a d-degenerate graph as the property is hereditary.

Let W be the set of nodes in G of degree at most $4d-2$. From Lemma 3.2 we know that $|W| \geq n/2$. Now let us call α_W the size of a maximum independent set of the graph induced by W. From Proposition 3.1 we have that

$$\alpha_W \geq \frac{|W|}{4d-2} \geq \frac{n}{2(4d-2)}.$$

Then, from Lemma 3.12 we know that after two time-steps, the expected number of stabilized nodes is:

$$\frac{\alpha_W}{4^{4d-1}} \geq \frac{n}{2(4d-2)4^{4d-1}}.$$

Let us denote $c(d) = 2(4d-2)4^{4d-1}$. The previous bound implies that after

$$T \geq \frac{2}{\log(c(d)) - \log(c(d)-1)} \cdot \log(n)$$

time-steps, the expected number of non-stabilized nodes is at most $1/n$. By the Markov inequality, we deduce that on time-step T the probability that all nodes are stabilized is at least $1 - 1/n$. We conclude that, with high probability, on $\mathcal{O}(\log n)$ time-steps all nodes are stabilized. □

Chapter 4

Conclusions and Future Work

4.1 Density Classification

In its original formulation, the density classification problem is stated for one-dimensional cellular automata. Since this problem has no solution, various efforts have been carried out in order to try to solve it perfectly in other formulations. In [7] the problem was stated in terms of the binary heat equation, with continuous states between 0 and 1, and shown that it could be solved on any d-dimensional regular grid with a sufficiently small value of the parameter σ, denoted *amplification factor*. In this thesis, we have proposed a generalization of this formulation to arbitrary graphs.

From the theoretical point of view, we have shown that any connected graph admits a solution for a sufficiently small amplification parameter. Later, we reported an experimental study of the dynamics. Our goal was to describe how the topology of a given graph G influences the convergence time and the effectiveness of the classification for fixed values of σ. Roughly, our results indicate that there is a noticeable difference between dense and sparse topologies with respect to the size of the amplification factor. In particular, connected acyclic graphs (trees) such as star graphs and Barabási-Albert graphs with parameter 1, require a smaller amplification factor compared to more dense topologies, such as Erdös-Renyi or regular graphs, in order to obtain comparable classification effectiveness. Since the convergence time of the dynamics is indirectly proportional to the size of σ, our results suggest that, from the computational complexity point of view, the density classification problem may be harder on sparse topologies than on dense topologies. The theoretical verification of this claim is a matter of future research.

4.2 Maximal Independent Set

We have presented a very simple dynamics that converges to configurations that represent a maximal independent set of the input graph on any initial configuration with probability 1. Our experimental results suggest that in average the convergence time of our dynamics is

$\mathcal{O}(\log n)$. Moreover, the dynamics on three or more states converges in $\mathcal{O}(\log n)$ steps with high probability, while in the case of two states, the convergence time is $\mathcal{O}(\log^2 n)$ with high probability.

The results given in Theorems 3.8 and 3.11 confirm the observations on graphs with constant independence number (such as complete graphs). Theorem 3.13 also confirms the observations of the 2-MIS-Dynamics on graphs of bounded degeneracy. Finally, Proposition 3.3 together with Theorems 3.8 and 3.11 implies poly-logarithmic convergence time for graphs of Erdös-Rényi in average, as well as with high probability.

We remark that the convergence time of our dynamics is not settled for general graphs. The existence of graph classes with large convergence time (on average or with high probability) is a possibility that we do not completely explore in this article. In this sense, we conjecture that our results can be improved in order to show a poly-logarithmic convergence time for every graph class.

From a more general perspective, we believe that there is an interesting research line related to the definition of simple finite-state dynamics that can work as strategies to efficiently compute to other graph structures, different that maximal independent sets, such as maximal matchings, minimal dominating sets, etc.

Bibliography

[1] Yehuda Afek, Noga Alon, Ziv Bar-Joseph, Alejandro Cornejo, Bernhard Haeupler, and Fabian Kuhn. Beeping a maximal independent set. *Distributed computing*, 26(4):195–208, 2013.

[2] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.

[3] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of algorithms*, 7(4):567–583, 1986.

[4] Leonid Barenboim, Michael Elkin, and Fabian Kuhn. Distributed (δ+1)-coloring in linear (in δ) time. *SIAM Journal on Computing*, 43(1):72–95, 2014.

[5] Béla Bollobás and Bela Bollobas. *Modern graph theory*, volume 184. Springer Science & Business Media, 1998.

[6] Béla Bollobás and Paul Erdös. Cliques in random graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 80, pages 419–427. Cambridge University Press, 1976.

[7] Raimundo Briceño, Pablo Moisset de Espanés, Axel Osses, and Iván Rapaport. Solving the density classification problem with a large diffusion and small amplification cellular automaton. *Physica D: Nonlinear Phenomena*, 261:70–80, 2013.

[8] Robert Stephen Cantrell and Chris Cosner. *Spatial ecology via reaction-diffusion equations*. John Wiley & Sons, 2004.

[9] Mathieu S Capcarrere, Moshe Sipper, and Marco Tomassini. Two-state, r= 1 cellular automaton that classifies density. *Physical review letters*, 77(24):4969, 1996.

[10] Yair Caro. New results on the independence number. Technical report, Technical Report, Tel-Aviv University, 1979.

[11] Shui-Nee Chow. Lattice dynamical systems. In *Dynamical systems*, pages 1–102. Springer, 2003.

[12] Pedro PB de Oliveira. On density determination with cellular automata: Results, constructions and directions. *Journal of Cellular Automata*, 9, 2014.

[13] Pedro PB de Oliveira, José C Bortot, and Gina MB Oliveira. The best currently known class of dynamically equivalent cellular automata rules for density classification. *Neurocomputing*, 70(1-3):35–43, 2006.

[14] Paula Gonzaga de Sá and Christian Maes. The gács-kurdyumov-levin automaton revisited. *Journal of Statistical Physics*, 67(3):507–522, 1992.

[15] Marcin Dembowski, Barbara Wolnik, Witold Bołt, Jan M Baetens, and Bernard De Baets. Two-dimensional affine continuous cellular automata solving the relaxed density classification problem. *Journal of Cellular Automata*, 14, 2019.

[16] Charles M Elliott and Zheng Songmu. On the cahn-hilliard equation. *Archive for Rational Mechanics and Analysis*, 96(4):339–357, 1986.

[17] Yuval Emek and Roger Wattenhofer. Stone age distributed computing. In *Proceedings of the 2013 ACM symposium on Principles of distributed computing*, pages 137–146, 2013.

[18] Nazim Fatès. Stochastic cellular automata solve the density classification problem with an arbitrary precision. In *Symposium on Theoretical Aspects of Computer Science-STACS2011*, volume 9, pages 284–295. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2011.

[19] Henryk Fukś. Solution of the density classification problem with two cellular automata rules. *Physical Review E*, 55(3):R2081, 1997.

[20] Péter Gács, Georgy L Kurdyumov, and Leonid Anatolevich Levin. One-dimensional uniform arrays that wash out finite islands. *Problemy Peredachi Informatsii*, 14(3):92–96, 1978.

[21] Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 270–277. SIAM, 2016.

[22] Mohsen Ghaffari, Christoph Grunau, and Václav Rozhoň. Improved deterministic network decomposition. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2904–2923. SIAM, 2021.

[23] George Giakkoupis. Simple efficient distributed processes on graphs, 2022. Keynote talk of the 29th International Colloquium of Structural Information and Communication Complexity, SIROCCO 2022, Paderborn, Germany, June 28th, 2022.

[24] Eric Goles, Laura Leal, Pedro Montealegre, Iván Rapaport, and Martín Ríos-Wilson. Distributed maximal independent set computation driven by finite-state dynamics. *International Journal of Parallel, Emergent and Distributed Systems*, 38:85–97, 2023.

[25] Sandra M Hedetniemi, Stephen T Hedetniemi, David P Jacobs, and Pradip K Srimani. Self-stabilizing algorithms for minimal dominating sets and maximal independent sets. *Computers & Mathematics with Applications*, 46(5-6):805–811, 2003.

[26] Hugues Juille and Jordan B Pollack. Coevolving the "ideal" trainer: Application to the discovery of cellular automata rules. In *University of Wisconsin*. Citeseer, 1998.

[27] Richard M Karp and Avi Wigderson. A fast parallel algorithm for the maximal independent set problem. *Journal of the ACM (JACM)*, 32(4):762–773, 1985.

[28] James Keener and James Sneyd. *Mathematical physiology: II: Systems physiology*. Springer, 2009.

[29] Fabian Kuhn, Thomas Moscibroda, and Rogert Wattenhofer. What cannot be computed locally! In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 300–309, 2004.

[30] Mark Land and Richard K Belew. No perfect two-state cellular automata for density classification exists. *Physical review letters*, 74(25):5148, 1995.

[31] Laura Leal, Pedro Montealegre, Axel Osses, and Iván Rapaport. A large diffusion and small amplification dynamics for density classification on graphs. *International Journal of Modern Physics C*, 34(05), 2023.

[32] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on computing*, 21(1):193–201, 1992.

[33] Michael Luby. A simple parallel algorithm for the maximal independent set problem. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 1–10, 1985.

[34] Claudio LM Martins and Pedro PB de Oliveira. Evolving sequential combinations of elementary cellular automata rules. In *European Conference on Artificial Life*, pages 461–470. Springer, 2005.

[35] Melanie Mitchell, James P Crutchfield, and Peter T Hraber. Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D: Nonlinear Phenomena*, 75(1-3):361–391, 1994.

[36] Norman H Packard. Adaptation toward the edge of chaos. *Dynamic patterns in complex systems*, 212:293–301, 1988.

[37] Stanislaw Ulam et al. Random processes and transformations. In *Proceedings of the International Congress on Mathematics*, volume 2, pages 264–275. Citeseer, 1952.

[38] Sébastien Verel, Philippe Collard, Marco Tomassini, and Leonardo Vanneschi. Fitness landscape of the cellular automata majority problem: view from the "olympus". *Theoretical Computer Science*, 378(1):54–77, 2007.

[39] John Von Neumann, Arthur W Burks, et al. Theory of self-reproducing automata. *IEEE Transactions on Neural Networks*, 5(1):3–14, 1966.

[40] Barbara Wolnik, Marcin Dembowski, Witold Bołt, Jan M Baetens, and Bernard De Baets. The density classification problem in the context of continuous cellular automata. In *International Conference on Cellular Automata*, pages 79–87. Springer, 2016.

[41] Dietmar Wolz and Pedro PB De Oliveira. Very effective evolutionary techniques for searching cellular automata rule spaces. *Journal of Cellular Automata*, 3(4), 2008.

[42] Jack Xin. Front propagation in heterogeneous media. *SIAM review*, 42(2):161–230, 2000.