

DISS. ETH NO. 23124

Society in Graphs

A thesis submitted to attain the degree of
DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

BARBARA KELLER

MSc, ETH Zürich, Switzerland

born on 22.04.1982

citizen of
Switzerland

accepted on the recommendation of
Prof. Dr. Roger Wattenhofer, examiner
Prof. Dr. Eric Goles, co-examiner
2015

Abstract

Individuals are seldom completely independent and should be considered as part of their social surroundings. Their actions both influence their social network and are influenced by the network. This phenomenon can be seen in online social networks formed on online platforms as well as in offline social networks formed, for example, by a group of scholars. In order to understand the fundamental effects of influence in social systems, we model social networks by the means of a well understood mathematical model, a graph.

The first part of this dissertation focuses on the impact of gender on the formation of a network of PhD students and their supervisors. We find evidence for homophilic behavior as well as for the existence of a *glass ceiling* in a network from the data of co-authorship. From over 1.3 millions of authors of scientific articles in computer science, we extract the student-supervisor relationship graph and analyze its properties.

Furthermore, we introduce mathematical formulations for the occurrence of a glass ceiling and an influence inequality in a network. We establish a network forming process integrating three observed characteristics of this network, namely a smaller entry rate for women, preferential attachment, and homophilic behavior. We prove that these three conditions are sufficient to produce a glass ceiling in the network. We also show, that if one of these three characteristic is missing, the glass ceiling according to the mathematical definition does not occur.

In the second part of this dissertation we examine how opinions evolve in networks. Every node in the network has an initial opinion and the nodes can observe the opinions of their neighbors. We assume a simplistic setting where the nodes are influenced by the opinions of their neighbors and always change their opinion to the opinion of the majority of their neighbors. We study several different variations of this model and investigate how long the system takes to reach a stable state. For asynchronous networks we find unweighted graphs which take $\Omega(n^2)$ steps until convergence. For unweighted synchronous networks we show graphs with a convergence time of $\tilde{\Omega}(n^2)$, where the $\tilde{\Omega}(\cdot)$ notation hides polylogarithmic factors. Additionally we show that allowing the influence to be weighted increases the convergence time dramatically to $\Omega(2^n)$.

Zusammenfassung

Menschen sind selten komplett unabhängig und sollten als Teil ihrer sozialen Umgebung betrachtet werden. Ihre Handlungen beeinflussen einerseits Ihre soziale Umgebung und sind andererseits selbst durch das soziale Netz der betreffenden Person beeinflusst. Man findet dieses Phänomen sowohl in sozialen Netzwerken welche auf online Plattformen geformt werden, als auch in herkömmlichen sozialen Netzwerken, wie man sie zum Beispiel unter Schülern findet. Um die fundamentalen Effekte von Einfluss in sozialen Systemen besser zu verstehen, modellieren wir soziale Netzwerke mit einem in der Mathematik gut bekannten Werkzeug, dem Graph.

Der erste Teil dieser Dissertation beschäftigt sich hauptsächlich mit den Auswirkungen welches das Geschlecht von Studenten auf die Bildung des Netzwerkes von Doktoranden/Doktorandinnen und ihrer Doktorväter/Doktormütter hat. In einem Netzwerk, aus Daten über Koautorenschaft, zeigen wir sowohl Hinweise für homophiles Verhalten als auch für die Existenz einer *gläsernen Decke*. Aus diesem Datenset, welches über 1.3 Millionen Autoren von wissenschaftlichen Artikeln aus dem Bereich Informatik enthält, haben wir einen Graph extrahiert, welcher die Studenten-Leiter Beziehung darstellt und studieren dessen Eigenschaften.

Ausserdem führen wir mathematische Formulierungen für die Erscheinung einer gläsernen Decke und einer ungleichen Verteilung von Einfluss in einem Netzwerk ein. Wir konstruieren einen Netzwerkbildungsprozess welcher drei beobachtete Eigenschaften, nämlich weniger Frauen, bevorzugte Anschliessung und homophiles Verhalten integriert. Zusätzlich beweisen wir, dass diese drei Eigenschaften ausreichend sind um eine gläserne Decke in einem Netzwerk hervorzurufen. Wir zeigen auch, dass falls eine dieser drei Eigenschaften wegfällt, sich keine gläserne Decke nach mathematischer Definition bildet.

Im zweiten Teil dieser Dissertation untersuchen wir, wie sich Meinungen in Netzwerken entwickeln. Jeder Knoten im Netzwerk hat eine ursprüngliche Meinung und die Knoten können die Meinung ihrer Nachbarn sehen. Wir nehmen ein einfaches Modell an, in welchem die Knoten beeinflusst sind von den Meinungen ihrer Nachbarn und ihre Meinung immer an die Meinung der Mehrheit Ihrer Nachbarn anpassen. Wir ana-

lysieren verschiedene Varianten dieses Modells und erforschen wie lange ein solches System braucht, um einen stabilen Zustand zu erreichen. In asynchronen Netzwerken finden wir ungewichtete Graphen welche $\Omega(n^2)$ Schritte benötigen um sich zu stabilisieren. Für ungewichtete, synchrone Netzwerke zeigen wir Graphen welche eine Stabilisierungszeit von $\tilde{\Omega}(n^2)$ benötigen, wobei die $\tilde{\Omega}(\cdot)$ Notation polylogarithmische Faktoren versteckt. Zusätzlich zeigen wir, dass wenn der Einfluss zwischen den Knoten gewichtet sein kann, sich die Stabilisierungszeit tiefgreifend erhöht zu $\Omega(2^n)$.

Acknowledgements

Being a PhD student has its ups and downs and I want to thank all the people who contributed to the ups and helped me over the downs. First of all my gratitude goes to my supervisor Roger Wattenhofer, who made this thesis possible. I enjoyed his openness and interest in all kinds of research questions. No matter if the topic concerned movies or same sex dating, he supplied me with interesting discussions and insights. Even insane sounding ideas, such as throwing your smartphone into the air to take aerial pictures, were welcomed. Among the many things he taught me, I appreciate the skills and courage to give good talks the most.

Then, I would like to thank my co-referee Eric Goles who did not only provide me with an upper bound for a research problem that caused many sleepless nights, but who was also fun and charming.

Next, I would like to thank Chen Avin, Zvi Lotker, and David Peleg, my supervisors during my exchange in Israel, for giving me the great opportunity to work on an exciting topic and to get to know a different way of life. I fondly remember the many ideas and the enthusiasm of Zvi, the many skills and interests of Chen, and the impressive knowledge and kindness of David. A very special thank you goes to Yvonne-Anne Pignolet, who helped me at so many stages during my studies and my PhD. Both, actively with her support as well as a role model of a successful and wonderful woman in computer science

My time would have had way fewer ups, if it was not for the DISCO members. I want to thank Georg Bachmeier for not drinking my whiskey away, Pascal Bissig for combining disgusting jokes with a cute interest in babies, Philipp Brandes for letting us use his compass on the road trip to Finland, Sebastian Brandt for his awesome hair (which is most likely the secret behind his genius), Friederike Bruetsch for helping to organize the KING summer school, Christian Decker for letting me use his crawling skills and his huge amount of memory, Raphael Eidenbenz for his encouraging words, Yuval Emek for his awesome teachings in spectral graph theory, Klaus-Tycho Förster for being always available at the office in any case of emergency, Silvio Frischknecht for his barefeet contributions to the long lost lower bound, Monica Fricker for her kind help in administration, Beat Futterknecht for solving every possible problem with a single phone

call and always lending a helping hand, Benny Gächter for solving all my soft- and hardware problems, Stephan Holzer for his point of view, Michael König for his awesomely mean game of super mario, Michael Kuhn for his contributions to the discussion at the Christmas dinner, Tobias Langner for his endless help with Latex and his shielding abilities, Tanja Lantz for her efficient help on administrative tasks, Christoph Lenzen for his contributions to the DISCO stories and publication list, Remo Meier for his silence, Laura Peer for being my most independent student, Johannes Schneider for his carefree attitude, Jochen Seidel for sharing cake and for many yes, Jasmin Smula for being a great office mate and spending submission nights with cute pink bunnies, Philipp Sommer for keeping the hardware in the group, David Stolz for being stolz, and Samuel Welten for defying gravity at work as well as in leisure time and letting me be part of both.

One of the perks as a PhD student is to work with bright students on interesting projects. I would like to thank all my students for their efforts and contributions and for the many successful projects as well as the interesting conversations: Nelly Afonso, Cyril Arnould, Anton Beitler, Hermann Blum, Dominik Blunschy, Gino Brunner, Conrad Burchert, Sivaranjini Chithambaram, Alexander Dietmüller, Kaspar Etter, Dominik Fankhauser, Markus Frei, Michael Giger, Lorenz Koestler, Jonas Luder, Marc Müller, Steffen Schmidt, Raphael Seebacher, and Philippe von Bergen.

I like to thank all my further collaborators: Orr Fischer, Juho Hirvonen, Tuomo Lempiäinen, Claire Mathieu, Joel Rybicki, and Jukka Suomela. I also want to thank Pierre Fraigniaud, Mauro Maggioni, Johan Paratte, Boaz Patt-Shamir, Nathanaël Perraudin, Jochen Seidel, and Pierre Vandergheynst for making our summer school "Key Insights on Networks and Graphs" a success. I would also like to thank my previous supervisor Stefan Wolf for the philosophical insights into randomness as well as for the puzzle solving night in the Safari bar.

Pandatar Matkustus and Panda Metaiel deserve a great part of my gratitude. They taught me how to travel properly and the value of coffee, not to mention their countless contributions to my research and talks.

For varying reasons I am deeply grateful to Larina Dettling, Ina Fischlin, Remo Gisi, Dominique im Obersteg, Adrian Merkle, Jeannine Pfeiffer, Chantal Reinhart, Andreas Ruedlinger, Tobias Schneebeli, Stephan

Wälti, Tamara Weilenmann, and Anita Weiss.

Furthermore, I want to thank my parents Lisa and Hansueli as well as my sisters Regula and Kathrin for all their support and sacrifices.

Finally, I am beholden to my husband candidate (and colleague) Jara Uitto for being a great partner in crime and for showing me that being embarrassing and being embarrassed does not need to go hand in hand. I am grateful to our little man Jarno for being such an awesome human being and I am looking forward to the μ 's to come.

Contents

1	Introduction	1
2	The Formation of a Social Network	5
2.1	Related Work	8
2.2	The Model and Definitions	10
2.3	Theoretical Results	17
2.4	Empirical Results	43
2.5	Conclusion	50
3	Social Network Dynamics	53
3.1	Related Work	55
3.2	Model	57
3.3	Sequential Influence Network	60
3.4	Synchronous Influence Network	65
3.5	Friends and Fiends	91
3.6	Weighted Influence Network	91
3.7	Asymmetric Influence Network	100
4	Conclusion	105

1

Introduction

Society is an interesting and mysterious construct exhibiting many funny, tragic, obvious and fascinating phenomena. Already the ancient Greeks were interested in studying the behavior of humankind. Since these times, when the terms society and state were used interchangeably, the research area of sociology has come a long way. Many different subareas such as sociology of knowledge, military sociology and gender studies, to just name a few, evolved. The methods used to investigate various aspects of societies also became more and more elaborate. Nowadays, the methods of sociology can be roughly divided into qualitative and quantitative methods. Qualitative methods traditionally include questionnaires with free form answers, interviews and first hand observations. These methods are suitable to analyze the motivation and (hidden) reasons for a certain behavior. In general, they are often used to deliver a descriptive analysis. Quantitative methods, on the other side, measure characteristics and

behavior in numbers and therefore provide a good basis for statistical arguments and classifications. This makes it easier to generalize statements and to compare them among different societies.

Clearly, both methods deserve researchers' attention and in particular thrive when their virtues are combined. While quantitative data can show particularities that should be investigated further, numbers rarely give a reason or an explanation. They, however, can help to understand what the important question is and where qualitative methods can assist to find explanations later on. Qualitative and quantitative methods also have a lot in common. Data is in the center of both approaches: gathering and analyzing it as well as drawing conclusions from it is the main goal. This is where the much younger research area of computer science comes into play. In computer science, data analysis is a well-studied subfield. Appropriately, sociology and computer science started flirting with each other, resulting in the research area of computational sociology.

For the first time in the history of mankind, we have the means and the data to analyze sociological phenomena in large scale networks with large scale data. As computers are designed to efficiently process masses of numbers, the focus lies on quantitative data. Today, we are able to scan gigantic text corpora for interesting patterns, which would take more than a life time to process manually. Also the available data grew enormously. Social online networks such as Facebook, Twitter and Wikipedia provide plenty of research material. On the one hand, they serve as data pool to investigate the online behavior of people. Sociologists study for example how people manage their privacy in social online networks or how they treat males and females differently in online games. On the other hand, the newly gained data is used to confirm long standing social theories in a larger environment, such as which factors favor segregation in a society or how people tend to (often subconsciously) manage their relationships by preferring to associate with friends of friends rather than with friends' friends. Computers can also contribute to facilitating qualitative methods. They can help to store data efficiently, preprocess transcripts and ease the input methods for qualitative data. With current technology, it is even possible to collect qualitative data more conveniently. Smartphones, for example, can be used to gather information about who is speaking in a meeting and already annotate the collected data accordingly.

Graph theory, a tool often used in computer science, proved itself particularly powerful in computational sociology. People are represented as nodes and their bonds are modeled as edges. Such a bond can represent friendship, collaboration, love interest, hate or whatever funny relationship scientists come up with. This, in turn, means that society is nothing but a graph (coincidentally the title of this thesis)!

In this work, we first shed light on a network forming process. As a female PhD student in the male dominated field of computer science, the most interesting network to study is the network of PhD students and their supervisors. Here, we focused particularly on the influence of gender on the formation of the network and the resulting consequences for the network. We modeled this process mathematically as a preferential attachment process, where the gender affects the attachment probabilities. During the formation process, the joining nodes show homophilic behavior and thereby influence the network structure. We rigorously analyze networks resulting from this process, and find that these networks exhibit a so-called *glass ceiling*, an invisible barrier preventing women from moving up the ladder. To evaluate how well our model captures reality, we analyze data sets of supervisors and their PhD students. Our biggest data set is extracted from a large publication data set on computer science, in which we reveal evidence for homophilic behavior as well as a glass ceiling.

In the second part of this work, we investigate how opinions evolve in existing networks if nodes try to influence each other's opinions. We consider a simple world with two alternative opinions spread over a network of opportunistic nodes. Each node adopts the opinion of the majority in its neighborhood. This leads to a process where the opinions of the nodes change until the network reaches a certain stable state. We analyze the convergence time of various kinds of influence networks, such as synchronous and asynchronous, as well as weighted, symmetric and asymmetric networks. We show that depending on the model, the convergence behavior of the network can change drastically. Interestingly, the simplest model where each edge represents an influence of weight 1, requires the most challenging analysis.

2

Homophily and the Glass Ceiling Effect in Social Networks

Attaining *equality of opportunity* is a fundamental value in democratic societies, therefore existing inequalities present us with a major concern. A particularly sore example is that many highly-qualified women and members of minority groups are unable to realize their full potential in society (and specifically in the workforce) due to a phenomenon commonly referred to as the *glass ceiling*, a powerful visual image for an invisible barrier blocking women and minorities from advancing past middle management levels [31]. A concern raised in a US Federal commission report [27]:

The “glass ceiling”... is the unseen, yet unbreakable barrier that keeps minorities and women from rising to the upper rungs of the corporate ladder, regardless of their qualifications or achievements.

The existence of the glass ceiling effect is well documented [16, 25, 60]. In academia, for example, gender disparities have been observed in the number of professors [67], earnings [22, 67, 82] funding [59] and patents [18]. A recent study [53] analyzed gender differences in research output, research impact and collaborations based on Thomson Reuters Web of Science databases. When prominent author positions were analyzed by sole authorship, first-authorship and last-authorship, it was discovered that papers with women in those leading roles were less frequently cited. The question we focus on in this article concerns the causes of this phenomenon. What are the invisible mechanisms that combine to create the glass ceiling effect, and in particular, what is the role of the social network in creating this effect? Many papers discuss possible causes of the glass ceiling effect and potential solutions to it, e.g., [17, 24, 51, 57], but to the best of our knowledge, this work is the first attempt to define the glass ceiling mathematically, study it in the context of the social network structure, and to propose a *mathematical model* capturing this phenomenon.

In order to talk about the glass ceiling we have to agree on a measure of success in a social network. Following the traditional approach that sees network edges as the “social capital” of the network, we define successful members of a social network to be high degree vertices, namely, the vertices that maintain a large number of connections, corresponding to high influence. Based on this we propose formal definitions for glass ceiling effects as a first contribution. Note that it is not clear how to capture the nature of a delicate dynamic mechanism like the glass ceiling in a concise yet precise way. To represent the dynamic nature we examine sequences of networks and their behavior when the number of vertices grows.

We take into account the following three well-accepted observations on human behavior related to forming networks, namely (i) the “rich get richer” mechanism, (ii) minority-majority partition (slower growth rate of a minority group in the network), and (iii) homophily (affinity towards those similar to oneself). The main result of this chapter is that under these three simple and standard assumptions the glass ceiling effect naturally arises in social networks. Let us first briefly describe these three social phenomena.

The “rich get richer” mechanism. This mechanism describes and explains the process of wealth concentration. It follows the basic idea that newly

created wealth is distributed among members of society in proportion to the amount they have already amassed. In our setting, where the degree of the vertex captures its level of social wealth, this mechanism predicts that people may try to connect more often to people who already have many connections. In order to profit from their social wealth or because they are more visible in the network.

Minority-majority partition. Many social groups exhibit unequal proportions of certain characteristics of their group members, one of them being gender. Even though we focus in this work on gender, the results can be applied to other characteristics as well. Certain occupations, such as construction, law enforcement, politics and computer science, tend to have a higher proportion of men. For example, the ratio of women taking up studies in the computing discipline varies per year and region between 10% and 35% [9, 38, 78, 86]. Other professions, such as elementary school teaching, nursing, and office administration, are occupied by a higher proportion of women. In fact, it is difficult to find an occupation with a balanced ratio of genders (this also holds for many other social partitions, e.g., ones based on ethnicity or family background). This imbalance is the second phenomenon underlying our model.

Homophily. It is a well established social phenomenon that people tend to associate with others who are similar to themselves. Characteristics such as gender, ethnicity, age, class background and education influence the relationships among human beings [54] and similarities make communication and relationship formation easier.

Based on these phenomena we propose a model obtained by applying the classical preferential attachment model (see Barabasi and Albert [6]) to a bi-populated minority-majority network augmented with homophily. The resulting model is hereafter referred to as the *Biased Preferential Attachment Model*. We prove that networks generated by this model exhibit a glass ceiling structure.

As a running example serving to illustrate the issue, let us consider the social network of mentor-student relationships in academia. With time, new (male and female) PhD students arrive and join the network. Upon arrival, each student needs to select a mentor. Over time, graduated students may become mentors themselves and some mentors become more

successful than others (e.g., in terms of the number of students they advise). How can one determine that there is a glass ceiling effect in this network? And if such an effect exists, what are the roots of its emergence? Is it merely a result of the females being a minority in the network, or is it some sort of discriminatory process? To complement our theoretical analysis with data from real networks we study these questions on a mentor-student network derived from vast data from publications in computer science. Using the definitions and insights from the model we observe the three earlier mentioned social phenomena as well as the models predictions of their effects on the network. Namely we find evidence for homophily, minority-majority partition, the “rich get richer” mechanism and the existence of glass ceiling effects.

2.1 Related Work

Homophily in social networks

Different characteristics such as gender, ethnicities, age, class background and education influence the relationships human beings form with each other [54]. McPherson et al. [64] survey a variety of properties and how they lead to particular patterns in bonding. Gender-based homophily can already be observed in play patterns among children at school [61, 83]. Eder and Hallinan [21] discovered that young girls are more likely to resolve intransitivity of friendships by deleting friendship choices, while young boys are more likely to add them. Overall, children are significantly more likely to resolve intransitivity by deleting a cross-sex friendship than by adding another cross-sex friendship [87]. These results show that gender influences the formation of cliques and of larger evolving network structures. These trends displaying homophily and gender differences in resolving problems in the structure of relationships mean that boys and girls gravitate towards different social circles. As adults, homophilic behavior persists, and men still tend to have networks that are more homophilic than women do. This behavior is even more pronounced in areas where men form the majority and in relationships exchanging advice and based on respect, e.g., mentoring [11, 40, 41, 81]. A homophilic network evolution model was studied in [10]. In this model new nodes connect to other nodes in two phases. First they choose their neighbors with a bias

towards their own type (the model allows a positive as well as a negative bias). In a second phase they choose their neighbors unbiased from the neighbors of their biased neighbors. The authors show, that the second phase overcomes the bias in the first phase and if the second phase is unbiased, the network ends up in an integrated state. They illustrate their model with data on citations in physics journals.

Gender disparity in science and technology

Gender disparities have been observed in the number of professors [22,67], earnings [82], funding [59] and patenting [18]. A related aspect is the “productivity puzzle”: men are more successful when it comes to number of publications and name position in the author list [89], for reasons yet unclear. Some conjectures raised involve (unknown) biased perceptions related to pregnancy/child care [13]. E.g., it was observed in [67] that science faculty members of both sexes exhibit unconscious biases against women. Simulations showed that even small male-female differences in work performance ratings can lead to substantially lower promotion rates for women, resulting in proportionately fewer women than men at the top levels of the organization [63]. Gender differences in research output, research impact and collaborations was analyzed in a study based on Thomson Reuters Web of Science databases [53]. It was not only revealed that papers with women in prominent author positions (sole authorship, first-authorship and last-authorship) were cited less frequently but the authors also found that age plays an important role in collaborations, authorship position and citations. Thus many of the trends observed therein might be explained by the under-representation of women among the elders of science. In other words, fixing the “leaky pipeline” [85] is key for a more equal gender distribution in science.

Minority of women in Computer Science

In the computing discipline, the ratio of women taking up studies varies by year and region between 10% and 35% [9, 38, 78, 86] (except in Malaysia, where women form a narrow majority [70]). This under-representation has been investigated [28, 84, 90] and remedial strategies have been proposed [26, 76]. There is a positive feedback loop [52]: the lack of women leads to a strong male stereotype which drives away even more women. Thus the increase of the relative number of women in computer science is argued to be the best of the investigated strategies, up to a “critical

mass” of women. However, as pointed out by Etkowitz [23], even achieving a critical mass of 15% women might not guarantee that the effects of a critical mass come into play.

2.2 The Model and Definitions

2.2.1 Biased preferential attachment model

We start by proposing a simple bi-populated preferential attachment model. In a gist, our model is obtained by applying the classical preferential attachment model (see Barabasi and Albert [6]) to a bi-populated minority-majority network augmented with homophily. The resulting model is hereafter referred to as the *Biased Preferential Attachment Model*. Formally, for $r \leq 1/2$ and $0 \leq \rho \leq 1$ let $G(n, r, \rho)$ be a variant of the preferential attachment model in which r represents the relative arrival rate of the minority vertices (to which we will refer to as the red vertices). Note that the expected fraction of red vertices in the network converges to r as well, as the relative size of the initial population becomes smaller over time. We define ρ to represent the level to which homophily, incorporated by using rejection sampling, is expressed in the system: for $\rho = 1$ the system is uniform and exhibits no homophily, whereas for $\rho = 0$ the system is fully segregated, and all added edges connect vertex pairs of the same color.

Let us describe the model in more detail. Denote the social network at time t by $G_t = (V_t, E_t)$, where V_t and E_t , respectively, are the sets of vertices and edges in the network at time t , and let $\delta_t(v)$ denote the degree of vertex v at time t . The process starts with an arbitrary initial (connected) network G_0 in which each vertex has an arbitrary color, red or blue. (For simplicity we require that a minimal initial network consists of one blue and one red vertex connected by an edge, but this requirement can be removed if $\rho > 0$). This initial network evolves in time as follows. In every time step t a new vertex v enters the network. This vertex is red with probability r and blue with probability $1 - r$. On arrival, the vertex v chooses an existing vertex $w \in V_t$ to attach to with probability p proportional to w 's degree at time t , i.e., $\mathbb{P}[w \text{ is chosen}] = \delta_t(w) / \sum_{u \in V_t} \delta_t(u)$.

Next, if w 's color is the same as v 's color, then an edge is inserted between v and w ; if the colors differ, then the edge is inserted with probability ρ , and with probability $1 - \rho$ the selection is *rejected*, and the process of choosing a neighbor for v is restarted. This process is repeated until some edge $\{v, w\}$ has been inserted. Thus in each time step, one new vertex and one new edge are added to the existing graph.

Figure 2.1 presents four examples of parameters for our model in the case of a 300-vertex bi-populated social network. First, Figure 2.1(a) provides an example for the *minority & homophily* case with $r = 0.3$ and $\rho = 0.7$ so the red vertices are a strict minority in the network and there is some homophily in the edge selection. The next three sub-figures present special cases. Figure 2.1(b) illustrates the *no minority* case (equal-size populations, i.e., $r = 0.5$) with homophily ($\rho = 0.7$). Figure 2.1(c) considers the *no homophily* case ($\rho = 1$) with minority ($r = 0.3$). The last extreme case, shown in Figure 2.1(d), is *absolute homophily*, where $\rho = 0$, and the red vertices are still in the minority ($r = 0.3$). This case results in *fully segregated* societies, namely, societies where members connect *only* to members of their own color. In this extreme case, the society in effect splits into two separate networks, one for each of the two populations (except for the single edge connecting the initial red and blue vertices).

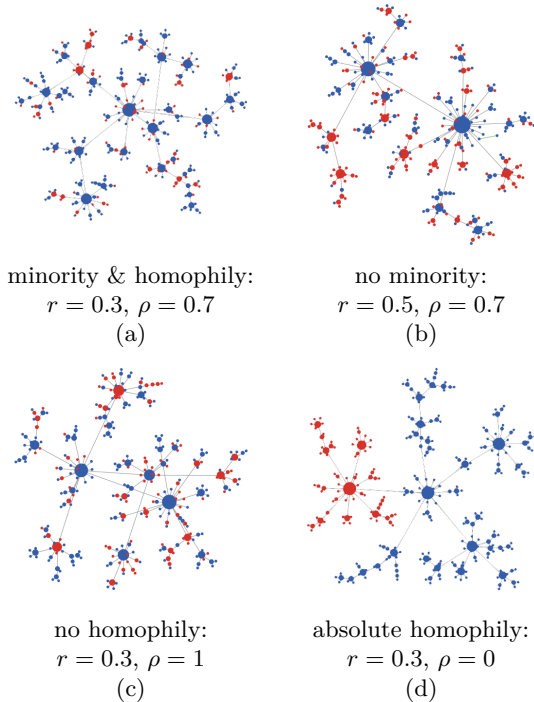


Figure 2.1: Examples of the Biased Preferential Attachment (BPA) model with various parameters. All examples depict a 300-vertex bi-populated network generated by our BPA model starting from a single edge connecting a blue and a red vertex (with vertex size proportional to its degree). (a) Minority & homophily: $r = 0.3$ (resulting in about 30% red vertices) and $\rho = 0.7$ (meaning that a new edge that connects red-blue vertices (i.e., a “mixed” edge) is accepted with probability 0.7 and otherwise rejected and sampled again, according to Preferential Attachment). (b) No minority & homophily: $r = 0.5$ and $\rho = 0.7$. (c) Minority & no-homophily: $r = 0.3$ and $\rho = 1$. (d) Minority & absolute homophily: $r = 0.3$ and $\rho = 0$ (indicating complete homophily in edge selection which results in two separate networks, one for the red vertices and the second for the blue vertices, plus a single initial connecting edge).

Consider as an example for our model the social network of mentor-student relationships in academia. With time, new PhD students arrive, but for some fields female students arrive at a lower rate than male students. Upon arrival, each student needs to select exactly one mentor, where the selection process is governed by the mechanisms of preferential attachment and homophily. Initially the student selects the mentor according to the rules of preferential attachment and then homophily takes its role, rejecting the selection with some probability if their gender differs, enforcing a re-selection. Over time, graduated students may become mentors and some mentors become more successful than others (in terms of the number of students they advise). Such a network exhibits a glass ceiling effect if, after a long enough time interval, the fraction of females among the most successful mentors tends to zero.

We would like to emphasize that the homophily effect sounds minor and “seemingly harmless”, in two ways. First, it is “symmetric”, i.e., it applies both to male students with respect to female mentors and to female students with respect to male mentors. Second, it does not adversely affect the student, in the sense that the student always gets admitted in our model. The only tiny (but ominous) sign for the potential danger of this homophilic effect is that it does affect the professor: a male professor who rejects (or is rejected by) some fraction of the female candidates risks little, whereas a female professor who rejects (or is rejected by) some fraction of the male candidates will eventually have fewer students overall, since most of the applicants are male. In fact, as we show later on, this homophily-based consequence will only impact her if her future potential students use preferential attachment to select their mentors.

2.2.2 Influence inequality and glass ceiling

Our second contribution is to propose formal definitions of the glass ceiling effect in social networks. Consider a *bi-populated* network $G(n)$ consisting of m edges and n nodes of two types, the group \mathbf{R} and the group \mathbf{B} . We assume that the network size n tends to infinity with time. Let $n(\mathbf{R})$ and $n(\mathbf{B})$, respectively, denote the number of red and blue nodes, where $n(\mathbf{R}) + n(\mathbf{B}) = n$. The red nodes are assumed to be a minority in the social network, i.e., denoting the percentage of red nodes in the network by r , we

assume $0 \leq r < \frac{1}{2}$. For a node v in G , let $\delta(v)$ denote its degree. Let $d(\mathbf{R})$ and $d(\mathbf{B})$ denote the sum of degrees of the red and blue nodes, respectively, where $d(\mathbf{R}) + d(\mathbf{B}) = 2m$. As the degree of a node corresponds to its power, the sum of the degrees of a certain kind of nodes represents the power of this group in the network. Let $\text{top}_k(\mathbf{R})$ (respectively, $\text{top}_k(\mathbf{B})$) denote the number of red (resp., blue) nodes that have degree *at least* k in G . When $G(n)$ is a random graph, we replace variables by their expectations in the definitions below, e.g., we use $\mathbb{E}[n(\mathbf{R})]$, $\mathbb{E}[d(\mathbf{R})]$, and $\mathbb{E}[\text{top}_k(\mathbf{R})]$. Next we provide formal definitions for the social phenomena discussed in the introduction. *Influence inequality* for the minority is defined in the following way.

Definition 2.1 (Influence inequality). A graph sequence $G(n)$ exhibits a *influence inequality* effect for the red nodes if the average power of a red node is lower than that of a blue (or a random) node, i.e., there exists a constant $c < 1$ such that

$$\lim_{n \rightarrow \infty} \frac{\frac{1}{n(\mathbf{R})} \sum_{v \in \mathbf{R}} \delta(v)}{\frac{1}{n(\mathbf{B})} \sum_{v \in \mathbf{B}} \delta(v)} = \frac{d(\mathbf{R})/n(\mathbf{R})}{d(\mathbf{B})/n(\mathbf{B})} \leq c. \quad (2.1)$$

The definition of the glass ceiling effect is more complex. We interpret the most powerful positions as those held by the highest degree nodes, and offer two alternative definitions. The first tries to capture the informal, “dictionary” definition, which describes a decreasing fraction of women among higher degree nodes, i.e., in the *tail* of the graph degree sequence. Formally:

Definition 2.2 (Tail glass ceiling). A graph sequence $G(n)$ exhibits a *tail glass ceiling* effect for the red nodes if there exists an increasing function $k(n)$ (for short k) such that $\lim_{n \rightarrow \infty} \text{top}_k(\mathbf{B}) = \infty$ and

$$\lim_{n \rightarrow \infty} \frac{\text{top}_k(\mathbf{R})}{\text{top}_k(\mathbf{B})} = 0.$$

The second definition considers a more traditional, distribution-oriented measure, the second moment of the two degree sequences. Formally:

Definition 2.3 (Moment glass ceiling). A graph sequence $G(n)$ exhibits a *moment glass ceiling* g for the red nodes where

$$g = \lim_{n \rightarrow \infty} \frac{\frac{1}{n(\mathbb{R})} \sum_{v \in \mathbb{R}} \delta(v)^2}{\frac{1}{n(\mathbb{B})} \sum_{v \in \mathbb{B}} \delta(v)^2}.$$

When $g = 0$, we say that $G(n)$ has a *strong* glass ceiling effect. The intuition behind this definition is that a larger second moment (and assuming a similar average degree, i.e., no influence inequality) will result in a larger variance and therefore a significantly larger number of high degree nodes. As we show later, the above two definitions for the glass ceiling are independent, in the sense that neither of the effects implies the other.

Note that these definitions are very general and do not rely on any assumptions of the degree distribution. In particular it is not necessary for networks that exhibit a glass ceiling effect to follow a power law degree distribution.

In order to investigate if a network reflects homophilic behaviour, we need to test define a test. Testing for *homophily* in a bi-populated network is based on checking whether the number of *mixed* (i.e., red-blue) edges is significantly lower than to be expected if neighbors were picked randomly and independently of their color. Formally:

Definition 2.4 (Homophily Test). [20] A bi-populated social network exhibits homophily if the fraction of mixed edges is significantly less than $2r(1-r)$.

The above definition implicitly assumes that there is power equality between the colors and therefore is not always accurate. A more careful test should take the average degree of each gender into account.

Definition 2.5 (Normalized Homophily Test). A bi-populated social network exhibits homophily if the fraction of mixed edges is significantly less than $2 \frac{d(\mathbb{R})}{2m} \left(1 - \frac{d(\mathbb{R})}{2m}\right)$.

An illustration of these definitions can be found in Figure 2.2.




Figure 2.2: (a) An example bi-populated social network with blue and red populations of 6 and 4 vertices respectively. (b) The degree sequences of both populations (i.e., the sequence specifying for each vertex its degree in the network). Considering the *tail glass ceiling* definition, there are four blue vertices of degree greater or equal to 4, but only two such red vertices so $\text{top}_4(\mathbf{R})/\text{top}_4(\mathbf{B}) = 1/2$. For the *moment glass ceiling* definition, the second moment for the blue vertices is $\frac{1}{6}(8^2 + 7^2 + 5^2 + 4^2 + 3^2 + 3^2) = 28.6$, while for the red vertices it is $\frac{1}{4}(7^2 + 5^2 + 3^2 + 3^2) = 23$ and the ratio is $23/28.6$. To exhibit a glass ceiling, these ratios should converge to zero as the network size increases. The average degree of the blue vertices in the network is 5 while the average for the red vertices is 4.5. It is possible that this numbers remain (almost) the same while the network size increases and the network exhibits a glass ceiling. Regarding homophily, in a random network with the same population, i.e., 60% blue vertices and 40% red vertices, one expects to find 36% blue-blue edges, 16% red-red edges and 48% mixed edge. If we take the degree sequences into account we would expect to see 46.8% mixed edge. In the above example network we observe only about 33% mixed edge, which indicates the effect of homophily.

2.3 Theoretical Results

2.3.1 Influence inequality and glass ceiling

Our main theoretical result (Thm. 2.1) is that in the biased preferential attachment model, $G(n, r, \rho)$, the glass ceiling effect emerges naturally. Additionally, this process generates a *influence inequality*, an independent property that is weaker than the glass ceiling effect. Influence inequality describes the situation where the average degree of the minority is lower than that of the majority (although their members possess the same qualifications). Moreover, we also show (Thm. 2.2) that all three ingredients (unequal entry rate, homophily, preferential attachment) are necessary to generate what we call a *strong* glass ceiling effect, i.e., removing any one of them will prevent the appearance of a glass ceiling effect. One may suspect that the glass ceiling effect is in fact a byproduct of influence inequality or unequal qualifications; we show that this is not the case. Minorities can have a smaller average degree without suffering from a glass ceiling effect. We also note that our results are independent of the starting condition. Even if the network initially consisted entirely of vertices of one color, if a majority of the vertices being added are of the opposite color, then eventually the vertices that rise to the highest positions will be of the new color.

Theorem 2.1. *Let $0 < r < \frac{1}{2}$ and $0 < \rho < 1$. For $G(n, r, \rho)$ produced by the Biased Preferential Attachment Model the following holds:*

- (i) $G(n, r, \rho)$ exhibits influence inequality, and
- (ii) $G(n, r, \rho)$ exhibits both a tail and a strong glass ceiling effect.

Moreover, all three ingredients are necessary to generate a *strong* glass ceiling effect.

Theorem 2.2. *(i) $G(n, r, \rho)$ will not exhibit a glass ceiling effect in the following cases:*

- (a) *If the rate $r = \frac{1}{2}$ (no minority).*
- (b) *If $\rho = 1$ (no homophily)*
- (c) *If $\rho = 0$ (no heterophily).*

- (ii) $G(n, r, \rho)$ will not exhibit a strong glass ceiling effect if the attachment process is uniform rather than preferential, i.e., a new vertex at time t selects an existing vertex to attach to uniformly at random from all vertices present at time $t - 1$ (and for any value of r and ρ).

Let us graphically illustrate the above results. Figure 2.3 presents the degree distributions of both the red and blue populations (as well as of the entire population) for four 1,000,000-vertex networks with parameters identical to the examples in Figure 2.1. The plots clearly show (and we prove this formally) that in all cases the degree distribution of both populations follows a power law. (A subset W of vertices in a given network obeys a power law degree distribution if the fraction $P(k)$ of vertices of degree k in W behaves for large values of k as $P(k) \sim k^{-\beta}$ for parameter β .) All figures present (in log-log scale) the cumulative degree distributions, so a power law degree distribution corresponds to a straight line (we present the samples together with the best-fit line). Theorem 2.1 corresponds to Figure 2.3(a) with the *minority & homophily* settings of $0 < r < \frac{1}{2}$ and $0 < \rho < 1$. In this case (and only in this case), the power law exponents of the red and blue populations, $\beta(\mathbf{R})$ and $\beta(\mathbf{B})$ respectively, are *different*, where $\beta(\mathbf{R}) > \beta(\mathbf{B})$; we prove that this will eventually lead to both a tail and a strong glass ceiling effect for the red vertices. Theorem 2.2 corresponds to Figures 2.3(b) and 2.3(c). The figures show that in the case of *no minority* (i.e., $r = 0.5$) or *no homophily* (i.e., $\rho = 1$), both $\beta(\mathbf{R})$ and $\beta(\mathbf{B})$ are the same (in particular they are equal to 3 as in the classical preferential attachment model), and therefore there will be no glass ceiling effect. Figure 2.3(d) considers the last extreme case of *absolute homophily*. Perhaps surprisingly, in this case a glass ceiling effect also does not occur, as each sub-population forms an absolute majority in its own network (see again Figure 2.1(d)). The case of no preferential attachment (which does not lead to a glass ceiling) is more delicate and presented in Section 2.3.4.




Figure 2.3: Graphical illustrations of our formal claims concerning the glass ceiling effect in the Biased Preferential Attachment model. Each figure presents the degree distribution (on a log-log scale) of the red and blue populations from a 1,000,000-vertex network generated by the BPA model with the same parameters as the corresponding figure in Figure 2.1. (a) Minority & homophily: $r = 0.3$ and $\rho = 0.7$. Both populations exhibit a power law degree distribution but with different exponents. Since $\beta(\mathbf{R}) > \beta(\mathbf{B})$, there is a glass ceiling effect for the red vertices. (The “noise” on the right-hand side of the graph stems from the fact that there are much fewer samples at the high-end of the range.) (b) No minority & homophily: $r = 0.5$ and $\rho = 0.7$. Both populations exhibit a power law degree distribution with $\beta = 3$, which indicates no glass ceiling effect. (c) Minority & no-homophily: $r = 0.3$ and $\rho = 1$. Again, the distributions indicate no glass ceiling effect. (d) Minority & absolute homophily: $r = 0.3$ and $\rho = 0$. Again, the distributions indicate no glass ceiling effect.

Proof Overview of Theorem 2.1. The basic idea behind the proof of Theorem 2.1 is to show that both populations in $G(n, r, \rho)$ have a power law degree distribution but with different exponents. Once this is established, it is simple to derive the glass ceiling effect for the population with a higher exponent in the degree distribution. To study the degree distribution of the red (and similarly the blue) population, we define α_t to be the random variable denoting the ratio of the total degree of the red nodes (i.e., the sum of degrees of all red nodes) to sum of all degrees (i.e., twice the number of edges). We show that the expected value of α_t converges to a fixed ratio independently of how the network started. The proof of this part is based on tools from dynamic systems, giving us directly a proof for 2.1 Part (i). Basically, we show that there is only one fixed point for our system. However, determining the expectation of α_t is not sufficient for analyzing the degree distribution, and it is also necessary to bound the rate of convergence and the concentration of α_t around its expectation. We used Doob martingales for this part. Using the high concentration of the total degree, we were able to adapt standard techniques to prove the power law degree distribution. Next we give an overview of the proofs and the helping lemmas.

2.3.2 Proof of Theorem 2.1 Part (i)

An urn process. The biased preferential attachment model $G(n, r, \rho)$ process can also be interpreted as a Polya's urn process, where each edge in the graph corresponds to two balls, one for each endpoint, and the balls are colored by the color of the corresponding vertices. When a new (red or blue) ball y arrives, we choose an existing ball c from the urn uniformly at random; if c is of the same color as y , then we add to the urn both y and another ball of the same color as c ; otherwise (i.e., if c is of a different color), with probability ρ we still add to the urn both y and another ball of the same color as c , and with probability $1 - \rho$ we reject the choice of c and repeat choosing an existing ball c' from the urn uniformly at random. Note that in any case c is put back into the urn. To analyze influence inequality, there is no need to keep track of the degrees of individual vertices; the sum of the degrees of all vertices of \mathbf{R} is exactly the number of red balls in the urn.

Denote by $u_t(\mathbf{R})$ (respectively, $u_t(\mathbf{B})$) the number of red (resp., blue) balls present in the urn at time $t \geq 0$. Altogether, the number of balls at time t is $u_t = u_t(\mathbf{R}) + u_t(\mathbf{B})$. Initially, the system contains u_0 balls. Noting that exactly two balls join the system in each time step, we have $u_t = u_0 + 2t$. Note that while $u_t(\mathbf{R})$ and $u_t(\mathbf{B})$ are random variables, u_t is not. Denote by α_t the random variable equal to $u_t(\mathbf{R})/u_t$, the fraction of red balls in the system at time t .

Convergence of expectations. We first claim that the process of biased preferential attachment converges to a ratio of α red balls in the system. More formally, we claim that regardless of the starting condition, there exists a limit

$$\alpha = \lim_{n \rightarrow \infty} \mathbb{E}[\alpha_t]. \quad (2.2)$$

We prove our claim step by step and start with presenting a function $F(x)$. The function $F(\alpha_t)$ describes the expected percentage growth of red balls at time $t + 1$ given the ratio of red balls at time t .

Lemma 2.6. $\mathbb{E}[\alpha_{t+1}|\alpha_t] = \alpha_t + \frac{F(\alpha_t) - \alpha_t}{t + 1}$, where

$$F(x) = \left(1 - (1 - r) \frac{(1 - x)}{1 - x(1 - \rho)} + r \frac{x}{1 - (1 - x)(1 - \rho)} \right) / 2.$$

Proof. We start from an arbitrary ratio $\alpha_0 = u_0(\mathbf{R})/u_0$. Observe that given that the new vertex is blue, the probability p that it attaches to a blue vertex satisfies $p = (1 - \alpha_t) + \alpha_t(1 - \rho)p$, hence $p = (1 - \alpha_t)/(1 - \alpha_t(1 - \rho))$. Given that the new vertex is red, the probability p' that it attaches to a red vertex satisfies $p' = \alpha_t + (1 - \alpha_t)(1 - \rho)p'$, hence $p' = \alpha_t/(1 - (1 - \alpha_t)(1 - \rho))$. We know that in each step the sum of the degrees increases by 2 in total so $u_{t+1} = u_t + 2$ and if X_t is the random variable that denotes the number of new red balls at time t , we obtain:

$$X_{t+1} = \begin{cases} 0 & \text{with probability } (1 - r) \frac{(1 - \alpha_t)}{1 - \alpha_t(1 - \rho)}, \\ & \text{(a blue ball entered and chose a blue ball)} \\ 2 & \text{with probability } r \frac{\alpha_t}{1 - (1 - \alpha_t)(1 - \rho)}, \\ & \text{(a red ball entered and chose a red ball)} \\ 1 & \text{with the remaining probability,} \\ & \text{(a blue ball chose a red ball or vice versa)} \end{cases}$$




Figure 2.4: Function $F(x)$ with the parameters $r = 1/3$ and $\rho = 1/2$, with its only fixed point in the interval $[0, 1]$ and $y = x$. The arrows represent the iterations of applying $F(x)$ repetitively onto itself, converging to its fixed point.

and we have $u_t(\mathbf{R}) = \sum_0^t X_i$. We now define:

$$\begin{aligned} \mathbb{E}[u_{t+1}(\mathbf{R}) - u_t(\mathbf{R}) | \alpha_t] &= 1 - (1 - r) \frac{(1 - \alpha_t)}{1 - \alpha_t(1 - \rho)} + r \frac{\alpha_t}{1 - (1 - \alpha_t)(1 - \rho)} \\ &= 2F(\alpha_t). \end{aligned}$$

Substituting $u_{t+1}(\mathbf{R}) = 2(t + 1)\alpha_{t+1}$ and $u_t(\mathbf{R}) = 2t\alpha_t$ and rewriting yields the Lemma. \square

We now have a function for the expected value of α_{t+1} given α_t . To prove that a_t actually converges to α we have to analyze the function $F(x)$ in more detail. We prove the following properties of this function:

Lemma 2.7. (i) $F(x)$ is monotonically increasing.

(ii) $F(x)$ has exactly one fixed point, denoted α^* , in $[0, 1]$.

(iii) The image of the unit interval by $F(x)$ is contained in the unit interval:

$$F([0, 1]) = \left[\frac{r}{2}, \frac{1+r}{2} \right] \subset [0, 1]$$

(iv) If $x < \alpha^*$ then $x < F(x) < \alpha^*$ and if $x > \alpha^*$ then $x > F(x) > \alpha^*$.

(v) $\alpha^* < r$.

Proof. With a little bit of algebra, we can, and for some reason we prefer, to rewrite $F(x)$ as

$$F(x) = \frac{1}{2} \left(r + \frac{rx}{x + (1-x)\rho} + (1-r) \frac{x\rho}{x\rho + (1-x)} \right).$$

For the first property, using simple algebra we compute

$$\frac{\partial F(x)}{\partial x} = \frac{1}{2} \left(\frac{\rho - \rho r}{(1 + (\rho - 1)x)^2} + \frac{\rho r}{(\rho + x - \rho x)^2} \right) > 0$$

for each $x, r, \rho \in [0, 1]$.

For the second property, we define the function $G(x) = F(x) - x$. The roots of $G(x)$ correspond to the fixpoints of $F(x)$ so it is enough to show that $G(x)$ has exactly one real root in the interval $[0, 1]$. Using simple algebra it follows that $\frac{\partial G(x)}{\partial x} > 0$ for each $x, r, \rho \in [0, 1]$. Setting $G(x) = 0$ we get the following equation:

$$(2-4\rho+2\rho^2)x^3 + (5\rho-2-3\rho^2-2r+2r\rho)x^2 + (2r-2\rho-2r\rho+\rho^2)x + r\rho = 0. \quad (2.3)$$

We observe that $G(x) = -\infty$ when $x \rightarrow -\infty$ for each $\rho \in [0, 1]$ and that $G(0) = r\rho \geq 0$. Which induces that there are 1 or 3 roots in the interval $(-\infty, 0)$. Observing that $G(x) = \infty$ when $x \rightarrow \infty$ and evaluating $G(1) = rp - p \leq 0$ for each $\rho, r \in [0, 1]$ we see that there are 1 or 3 roots in the interval $(1, \infty)$. Knowing that $G(x)$ has exactly 3 roots concludes the claim that G has exactly one root in $[0, 1]$ which leads to the conclusion that the function $F(x)$ has exactly one fixed point, α^* , in $[0, 1]$.

The third property follows from the fact that the function $F(x)$ is strictly monotonically increasing and by evaluating the function $F(x)$ for the two extreme values $x = 0$, and $x = 1$.

The fourth property follows from the fact that the function is strictly monotonically increasing, that there is only one fixed point and that $F(x)$ maps $[0, 1]$ inside $[0, 1]$.

Finally, to show that $\alpha^* < r$, since we know that $F(x) - x$ is positive for $x < \alpha^*$ and negative for $x > \alpha^*$, it suffices to show that $F(r) < r$. This is equivalent to

$$r + \frac{r^2}{r + (1-r)\rho} + (1-r)\frac{r\rho}{r\rho + (1-r)} < 2r,$$

which is true for all $r < 1/2$. □

Now assume $\alpha_t < \alpha^*$. By Lemma 2.7, $\alpha_t < F(\alpha_t) < \alpha^*$, so by Lemma 2.6 we obtain $\alpha_t < \mathbb{E}[\alpha_{t+1} | \alpha_t] < \alpha^*$.

With that we have shown that the expected value of α_t does converge to the fixed point α^* of $F(x)$. Figure 2.4 shows an instance of $F(x)$ with the parameters $r = 1/3$ and $\rho = 1/2$. You can see its only fixed point in the interval $[0, 1]$ on the intersection of $F(x)$ with the line $y = x$. If the function $F(x)$ is applied repetitively it will converge to its fixed point, no matter if the initial value was larger or smaller than the fixed point. We still need to bound the rate of convergence of $F(x)$.

Auxiliary Lemma. In order to bound the rate of convergence we will need the following lemma. It states that for any $r/2 \leq \alpha^* \leq r$, the straight line joining the points $(0, r/3)$ and $(\alpha^*, F(\alpha^*))$ is below the function $F(x)$ for all $x \in [0, \alpha^*]$.

Lemma 2.8. *For all $r \in (0, 1/2)$, $\rho \in [0, 1]$, $r/2 \leq \alpha^* \leq r$, and $x \in [0, \alpha^*]$*

$$F(x) \geq \frac{r}{3} + x \frac{F(\alpha^*) - r/3}{\alpha^*}$$

Proof. We need to show that $\psi(x) \equiv F(x) - \frac{F(\alpha^*) - r/3}{\alpha^*} \cdot x - \frac{r}{3} \geq 0$. Note

that $\psi(x)$ can be rewritten as

$$\psi(x) = \frac{r}{6} \left(x \left(3 \left(w - \frac{1}{\alpha^*(-\rho) + \alpha^* + \rho} \right) - \frac{1}{\alpha^*} \right) + 1 \right) \quad (2.4)$$

$$+ \frac{\rho x}{2} \left(\frac{1}{\alpha^*(-\rho) + \alpha^* - 1} + \frac{1}{(\rho - 1)x + 1} \right) \quad (2.5)$$

$$\text{where } w = \rho \left(\frac{1}{\alpha^*(\rho - 1) + 1} + \frac{1}{-\rho x + x - 1} \right) + \frac{1}{\rho - \rho x + x}.$$

We arrange $\psi(x)$ as $\frac{N(x)}{D(x)}$, then the denominator can be written as:

$$D(x) = 6\alpha^*(\alpha^*\rho - \alpha^* + 1)(\alpha^*\rho - \alpha^* - \rho)(\rho x - x + 1)(-\rho + \rho x - x)$$

which is positive, so the sign of $\psi(x)$ is determined by the numerator $N(x)$. The advantage of working with the numerator is that it is a polynomial. Some calculations yield that

$$\begin{aligned} N(x) &= (x - \alpha^*)r\rho(1 + \alpha^*(\rho - 1))(-\rho + \alpha^*(\rho - 1)) \\ &\quad + (x - \alpha^*)r(\rho - 1)^2x^2 \left(2(\alpha^* - 1)\alpha^*(\rho^2 + \rho - 2) + \rho \right) \\ &\quad - (x - \alpha^*)r(\rho - 1)x2\alpha^{*2}(\rho^3 - 3\rho + 2) \\ &\quad - (x - \alpha^*)r(\rho - 1)x \left((\rho - 1)\rho - \alpha^*(\rho(2\rho + 3) - 3) + 4 \right) \\ &\quad - (x - \alpha^*)3\alpha^*(\rho - 1)\rho x(\alpha^*(\rho - 1) - \rho)((\rho - 1)x - \rho). \end{aligned}$$

Clearly $N(x)$ has a root at α^* , i.e., $N(\alpha^*) = 0$. Since the degree of $N(x)$ is 3 it follows that it has at most three real roots $\lambda_1 \leq \lambda_2 \leq \lambda_3$. Assume that $\lambda_2 = \alpha^*$. A simple calculation shows that the leading coefficient of the polynomial $N(x)$ is $\alpha^*\rho(-r)(\alpha^*(\rho - 1) + 1)(\alpha^*(\rho - 1) - \rho)$ and therefore it is positive. This implies that

$$\lim_{x \rightarrow -\infty} N(x) = -\infty$$

and that

$$\lim_{x \rightarrow \infty} N(x) = \infty.$$

Next we claim that $N[0] > 0$. Setting $x = 0$ in $N(x)$ we get that $N(0) = \alpha^* \rho(-r)(\alpha^*(\rho - 1) + 1)(\alpha^*(\rho - 1) - \rho)$, which is positive in the range of the variables r, α^*, ρ . This shows that some root of N is less than 0, i.e., $\lambda_1 < 0$.

Next we show that there is one root of N that is greater than $1/2$, i.e., $\lambda_3 > 1/2$. To do this, we consider $N(1/2)$ and show that it is negative, i.e., $N(1/2) < 0$. A simple calculation shows that

$$\begin{aligned} N(1/2) = & \left(\frac{1}{2} - \alpha^*\right)r\frac{1}{2}(1 - \rho)\left((\rho - 1)\rho - \alpha^*(\rho(2\rho + 3) - 3) + 4\right) \\ & + r\alpha^{*2}\left(\frac{1}{2} - \alpha^*\right)(1 - \rho)(\rho^3 - 3\rho + 2) \\ & + \left(\frac{1}{2} - \alpha^*\right)r\left(\frac{1}{4}(\rho - 1)^2(2(\alpha^* - 1)\alpha^*(\rho^2 + \rho - 2) + \rho)\right) \\ & + \left(\frac{1}{2} - \alpha^*\right)r\rho(\alpha^*(\rho - 1) + 1)(\alpha^*(\rho - 1) - \rho) \\ & - \left(\frac{1}{2} - \alpha^*\right)\frac{3}{2}\alpha^*\left(\frac{\rho - 1}{2} - \rho\right)(\rho - 1)\rho(\alpha^*(\rho - 1) - \rho). \end{aligned}$$

Viewing $N(0)$ as a function of r , we notice that $N(0)$ is a monotonically linear descending function of r and therefore one can assume that $r = \alpha^*$. In this case, i.e., $r = \alpha^*$, we get that

$$\begin{aligned} N(1/2) < & \left(\frac{1}{2} - \alpha^*\right)\alpha^*\frac{1}{2}(1 - \rho)\left((\rho - 1)\rho - \alpha^*(\rho(2\rho + 3) - 3) + 4\right) \\ & + \alpha^{*3}\left(\frac{1}{2} - \alpha^*\right)(1 - \rho)(\rho^3 - 3\rho + 2) \\ & + \left(\frac{1}{2} - \alpha^*\right)\alpha^*\left(\frac{1}{4}(\rho - 1)^2(2(\alpha^* - 1)\alpha^*(\rho^2 + \rho - 2) + \rho)\right) \\ & + \left(\frac{1}{2} - \alpha^*\right)\alpha^*\rho(\alpha^*(\rho - 1) + 1)(\alpha^*(\rho - 1) - \rho) \\ & - \left(\frac{1}{2} - \alpha^*\right)\frac{3}{2}\alpha^*\left(\frac{\rho - 1}{2} - \rho\right)(\rho - 1)\rho(\alpha^*(\rho - 1) - \rho). \end{aligned}$$

Now a simple calculation shows that this is a monotonically descending function of ρ . Setting $\rho = 0$, we get

$$N(1/2) < \left(\frac{1}{2} - \alpha^*\right)\alpha^*\left(\frac{1}{2}(4\alpha^{*2} - 4\alpha^*) + (1 - \alpha^*)\alpha^*\right),$$

which is less than 0 for all $0 < \alpha^* < 1/2$ so $\lambda_2 > 1/2$. This in turn implies that $\psi(x) \geq 0$ when $x \in [0, \alpha^*]$. \square

Now we can bound the rate of convergence of $F(x)$.

Lemma 2.9. $|\alpha^* - \mathbb{E}[\alpha_t]| = O(1/\sqrt[3]{t})$.

Proof. Assume that $\alpha_t < \alpha^*$ (the other case is similar).

$$\mathbb{E}[\alpha_{t+1}] = \mathbb{E}[\alpha_t] + \frac{\mathbb{E}[F(\alpha_t) - \alpha_t]}{t+1} \quad (2.6)$$

Let $\Delta = \frac{\alpha^* - r/3}{\alpha^*}$ and the line $L(x) = \Delta \cdot x + r/3$. Note that L shares the point $(\alpha^*, F(\alpha^*))$ with $F(x)$ but is strictly below $F(x)$ in the range $[0, \alpha^*]$ it (See Lemma 2.8). Thus

$$\mathbb{E}[F(\alpha_t)] \geq \mathbb{E}[\alpha^* - \epsilon_t \Delta],$$

where $\epsilon_t = \alpha^* - \alpha_t$.

Substituting into Equation (2.6), we get

$$\begin{aligned} \mathbb{E}[\alpha_{t+1}] &\geq \mathbb{E}[\alpha_t] + \frac{\mathbb{E}[\alpha^* - \epsilon_t \Delta - \alpha_t]}{t+1} \\ &= \mathbb{E}[\alpha_t] + \frac{\mathbb{E}[\epsilon_t](1 - \Delta)}{t+1} = \alpha^* - \mathbb{E}[\epsilon_t] + \frac{\mathbb{E}[\epsilon_t](1 - \Delta)}{t+1}, \end{aligned}$$

so the expected error at time $t+1$ is

$$\mathbb{E}[\epsilon_{t+1}] \leq \mathbb{E}[\epsilon_t] \left(1 - \frac{1 - \Delta}{t+1}\right).$$

Solving for $\mathbb{E}[\epsilon_t]$ we have

$$\begin{aligned} \mathbb{E}[\epsilon_t] &= \epsilon_0 \left(1 - \frac{1 - \Delta}{2}\right) \left(1 - \frac{1 - \Delta}{3}\right) \cdots \left(1 - \frac{1 - \Delta}{t}\right) \\ &= \epsilon_0 \exp\left(-\sum_{i=1}^t \frac{1 - \Delta}{i}\right) = O\left(\frac{\epsilon_0}{t^{1-\Delta}}\right), \end{aligned}$$

Note that since $r/2 \leq \alpha^* < r$ we have $\Delta < 2/3$, and so $\mathbb{E}[\epsilon_t] = O(1/\sqrt[3]{t})$. \square

With the previous steps we have proven both, the convergence of $\mathbb{E}[\alpha_t]$ to α^* as well as we have bound the convergence rate. We now investigate the fixed point α^* .

Theorem 2.3. *For any initial configuration, as t goes to infinity, the expected fraction of red balls in the urn, $\mathbb{E}[\alpha_t^*]$, converges to the unique α^* in $[0, 1]$ satisfying the equation*

$$2\alpha^* = 1 - (1-r) \frac{(1-\alpha^*)}{1-\alpha^*(1-\rho)} + r \frac{\alpha^*}{1-(1-\alpha^*)(1-\rho)}. \quad (2.7)$$

Reformulating Equation 2.7 shows that the limit α^* is the solution of the following cubic equation.

$$\begin{aligned} 0 = & (4\rho - 2\rho^2 - 2)(\alpha^*)^3 + (2 + 3\rho^2 - 5\rho + 2r - 2r\rho)(\alpha^*)^2 \\ & + (2\rho - 2r + 2r\rho - \rho^2)\alpha^* - r\rho \end{aligned} \quad (2.8)$$

Note that this α^* is independent of the initial values u_0 and α_0 of the system. From Lemma 2.7 Part (ii) we know that α^* exists and that α^* satisfies the Equations 2.7 and 2.8. Recall that in Equation 2.2 we defined the limit of $\mathbb{E}[\alpha_t]$ as α . Combining this with Lemma 2.9 we know that α_t converges to α^* . From that we can conclude that α exists and that $\alpha = \alpha^*$.

Having shown the independence of α of the initial configuration of the urn we make the following observation. The expected degree of a red node tends to $2\alpha/r$, which is strictly less than 2, the expected degree of a random node, because of Lemma 2.7 Part (v). This concludes the proof for Theorem 2.1 Part (i) and we have shown that $G(n, r, \rho)$ exhibits an influence inequality.

2.3.3 Proof of Theorem 2.1 Part (ii)

Concentration. To prove the glass ceiling effect we first bound the degree distribution. To do this we need to bound the rate by which $u_t(\mathbb{R})$ converge

to $\alpha \cdot t$. Let $X_i \in \{0, 1, 2\}$ be the number of new red balls in the system at time i . Note that $u_t(\mathbf{R}) = \sum_0^t X_i$. Let

$$\bar{X}_i = (X_1, X_2, \dots, X_i)$$

be a tuple that captures all random variables X_1, X_2, \dots, X_i and let

$$\Psi_i = \mathbb{E}_{X_{i+1}, X_{i+2}, \dots, X_t} \left[\sum_{j=0}^t X_j | \bar{X}_i \right].$$

Observe that $(\Psi_i)_i$ is a Doob Martingale [66], and note that we already know that $\Psi_0 = \mathbb{E} \left[\sum_{i=0}^t X_i \right] = \mathbb{E} \left[u_t(\mathbf{R}) \right]$.

Theorem 2.4 (Azuma's inequality [3]). *Let Ψ_t be a martingale such that for all i , almost surely $|\Psi_i - \Psi_{i-1}| < c_i$. Then for all positive t and all positive reals x ,*

$$\Pr(\Psi_t - \Psi_0 \geq x) \leq \exp \left(\frac{-x^2}{2 \sum_i c_i^2} \right).$$

Lemma 2.10. *Let $C_i = |\Psi_i - \Psi_{i-1}|$. Then $C_i = O(\sqrt{t/i})$.*

Proof. Observe that for $c = 0, 1, 2$,

$$\Psi_i - \Psi_{i-1} = \mathbb{E}_{X_{i+1}, \dots, X_t} \left[\left(\mathbb{E} \left(\sum_{j=i}^t X_j | X_i = c \right) - \mathbb{E}_{X_i} \left(\sum_{j=i}^t X_j \right) \right) | \bar{X}_i \right].$$

To bound $C_i = |\Psi_i - \Psi_{i-1}|$, since each additional ball creates an *independent* effect on its descendants (the red balls that are connected to it), we have: $\mathbb{E}(C_i) \leq 2z_i$, where

$$z_i = \mathbb{E} \left[a_{i,t} | b_i \right]$$

where $a_{i,t}$ = number of additional red balls at times $[i, t]$, and b_i = one additional red ball at time i . We have the recurrence: $z_t = 1$, and for $i < t$

$$z_i = 1 + \frac{\gamma}{2(i+1)} z_{i+1} + \frac{\gamma}{2(i+2)} z_{i+2} + \dots + \frac{\gamma}{2t} z_t$$

where $\gamma/(2i)$ is the probability of selecting a particular marked red ball at time i : we always have $\gamma \leq 1$, and γ depends on the homophily parameter ρ .

We apply some algebraic changes and let $y_i = z_i/2i$. It is easy to see that the recurrence becomes $y_t = 1/(2t)$, and $y_i = (1 + (2 + \gamma)/(2i))y_{i+1}$. Solving the recurrence yields

$$\begin{aligned} y_i &= \left(1 + \frac{2 + \gamma}{2i}\right) \left(1 + \frac{2 + \gamma}{2(i+1)}\right) \cdots \frac{1}{2t} \\ &= O\left(\frac{1}{2t} \exp\left(\sum_{j=i}^t \frac{2 + \gamma}{2j}\right)\right) \\ &= O\left(\frac{1}{2t} \left(\frac{t}{i}\right)^{\frac{2 + \gamma}{2}}\right). \end{aligned} \tag{2.9}$$

Since $\gamma \leq 1$, we obtain that $y_i = O((1/i)\sqrt{t/i})$ and we also obtain $|C_i| = O(z_i) = O(\sqrt{t/i})$. \square

Recall that $\alpha_t = u_t(\mathbf{R})/(2t)$. By Theorem 2.4 and Lemma 2.10 then

Lemma 2.11. $\Pr\left[|u_t(\mathbf{R}) - 2t\mathbb{E}(\alpha_t)| > O(2\sqrt{t} \log t)\right] \leq \frac{1}{t^4}$.

Proof. By Lemma 2.10

$$\sum_{i=1}^t C_i^2 \leq O\left(\sum_{i=1}^t \frac{t}{i}\right) \leq ct \log t.$$

We now use Theorem 2.4 for $x = \Theta(2\sqrt{ct} \log t)$, note that $\Psi_t = u_t(\mathbf{R})$ and $\Psi_0 = \mathbb{E}[u_t(\mathbf{R})]$, and obtain:

$$\Pr\left[|u_t(\mathbf{R}) - \mathbb{E}[u_t(\mathbf{R})]| > O(\sqrt{4t} \log t)\right] \leq 2 \exp\left(\frac{-4t \log^2 t}{t \log t}\right) = O\left(\frac{1}{t^4}\right).$$

\square

Combining Lemmas 2.9 and 2.11 yields:

Corollary 2.12.

$$\Pr \left[|\alpha_t - \alpha| > \max \left\{ \frac{2 \log t}{\sqrt{t}}, \frac{1}{\sqrt[3]{t}} \right\} \right] < \frac{1}{t^4}.$$

We have shown that the ratio of red balls converts fast enough to a certain α and is concentrated around the expected value of α . With this knowledge we can investigate the degree distribution of the red nodes.

Degree distribution

We investigate the degree distribution of the red and blue vertices in a graph generated by the biased preferential attachment process, following the analysis outline of [15] for the basic preferential attachment model.

Let $m_{k,t}(\mathbf{B})$ (resp., $m_{k,t}(\mathbf{R})$) denote the number of blue (resp., red) vertices of degree k at time t . For $\mathbf{x} \in \{\mathbf{R}, \mathbf{B}\}$, define

$$M_k(\mathbf{x}) = \lim_{t \rightarrow \infty} \frac{\mathbb{E}[m_{k,t}(\mathbf{x})]}{t}. \quad (2.10)$$

Theorem 2.5. *The expected degree distributions of the blue and red vertices follow a power law, namely, $M_k(\mathbf{B}) \propto k^{-\beta(\mathbf{B})}$ and $M_k(\mathbf{R}) \propto k^{-\beta(\mathbf{R})}$. If $0 < r < 1/2$ and $0 < \rho < 1$ then $\beta(\mathbf{R}) > 3 > \beta(\mathbf{B})$.*

As next we will prove Theorem 2.5. Note that $m_{0,0}(\mathbf{B}) = u_0(\mathbf{B})$. We derive a recurrence for $\mathbb{E}[m_{k,t}(\mathbf{B})]$. A blue vertex of degree k at time t could have arisen from three scenarios: (s1) at time $t - 1$ it was already a blue vertex of degree k and no edge was added to it at time t . (s2) at time $t - 1$ it was a blue vertex of degree $k - 1$ and an edge was added to it at time t . (s3) in the special case where $k = 1$, at time $t - 1$ it did not exist yet and it has arrived as a new blue vertex at time t . Thus letting \mathcal{F}_t be the history of the process up to time t , for any $k > 1$, the expectation of $m_{k,t+1}(\mathbf{B})$ conditioned on \mathcal{F}_t satisfies

$$\begin{aligned} \mathbb{E}[m_{k,t+1}(\mathbf{B})|\mathcal{F}_t] = & m_{k,t}(\mathbf{B}) \left(1 - \frac{ru_t(\mathbf{B})\rho \frac{k}{u_t(\mathbf{B})}}{u_t(\mathbf{R}) + u_t(\mathbf{B})\rho} - \frac{(1-r)u_t(\mathbf{B}) \frac{k}{u_t(\mathbf{B})}}{u_t(\mathbf{R})\rho + u_t(\mathbf{B})} \right) \\ & + m_{k-1,t}(\mathbf{B}) \left(\frac{ru_t(\mathbf{B})\rho \frac{k-1}{u_t(\mathbf{B})}}{u_t(\mathbf{R}) + u_t(\mathbf{B})\rho} + \frac{(1-r)u_t(\mathbf{B}) \frac{k-1}{u_t(\mathbf{B})}}{u_t(\mathbf{R})\rho + u_t(\mathbf{B})} \right). \end{aligned}$$

For $k = 1$ we similarly have

$$\mathbb{E}[m_{1,t+1}(\mathbf{B})|\mathcal{F}_t] = m_{1,t}(\mathbf{B}) \left(1 - \frac{\rho r}{u_t(\mathbf{R}) + u_t(\mathbf{B})\rho} - \frac{1-r}{u_t(\mathbf{R})\rho + u_t(\mathbf{B})} \right) + 1-r. \quad (2.11)$$

Recalling again that $\alpha_t = u_t(\mathbf{R})/(2t)$, the $\mathbb{E}[m_{k,t+1}(\mathbf{B})|\mathcal{F}_t]$ term from above can be rewritten as

$$\begin{aligned} & m_{k,t}(\mathbf{B}) \left(1 - \frac{r\rho k}{2t(\alpha_t + (1-\alpha_t)\rho)} - \frac{(1-r)k}{2t(\alpha_t\rho + (1-\alpha_t))} \right) \\ & + m_{k-1,t}(\mathbf{B}) \left(\frac{r\rho(k-1)}{2t(\alpha_t + (1-\alpha_t)\rho)} + \frac{(1-r)(k-1)}{2t(\alpha_t\rho + (1-\alpha_t))} \right). \end{aligned}$$

and for the case of $k = 1$, the term $\mathbb{E}[m_{1,t+1}(\mathbf{B})|\mathcal{F}_t]$ can be written as

$$m_{1,t}(\mathbf{B}) \left(1 - \frac{\rho r}{2t(\alpha_t + (1-\alpha_t)\rho)} - \frac{1-r}{2t(\alpha_t\rho + (1-\alpha_t))} \right) + (1-r).$$

This can be expressed as

$$\begin{aligned} \mathbb{E}[m_{k,t+1}(\mathbf{B})|\mathcal{F}_t] &= m_{k,t}(\mathbf{B}) \left(1 - A_t \frac{k}{t} \right) + m_{k-1,t}(\mathbf{B}) A_t \frac{k-1}{t}, \\ \mathbb{E}[m_{1,t+1}(\mathbf{B})|\mathcal{F}_t] &= m_{1,t}(\mathbf{B}) \left(1 - \frac{A_t}{t} \right) + (1-r), \end{aligned} \quad (2.12)$$

using the notation

$$A_t = \frac{r\rho}{2\alpha_t + 2(1-\alpha_t)\rho} + \frac{(1-r)}{2\alpha_t\rho + 2(1-\alpha_t)}.$$

Similar calculations hold for the red vertices. Substituting r by $(1 - r)$ and α_t by $(1 - \alpha_t)$ respectively leads to the corresponding factor for the red vertices. Note that A_t is a random variable so we next bound its divergence. Denoting by C_B the factor for the blue vertices and by C_R the factor for the red vertices. Let

$$C_B = \frac{r\rho}{2\alpha + 2(1 - \alpha)\rho} + \frac{(1 - r)}{2\alpha\rho + 2(1 - \alpha)} \quad (2.13)$$

$$C_R = \frac{(1 - r)\rho}{2\alpha\rho + 2(1 - \alpha)} + \frac{r}{2\alpha + 2(1 - \alpha)\rho}. \quad (2.14)$$

We have

Lemma 2.13. $\Pr \left[|A_t - C_B| > \max \left\{ \frac{2 \log t}{\sqrt{t}}, \frac{1}{\sqrt[3]{t}} \right\} \right] < \frac{1}{t^4}.$

To understand the degree distribution in our model we will make use the following lemma.

Lemma 2.14. [15] *Let $(a_t), (b_t), (c_t)$ be three sequences such that $a_{t+1} = (1 - \frac{b_t}{t})a_t + c_t$, $\lim_{t \rightarrow \infty} b_t = b > 0$ and $\lim_{t \rightarrow \infty} c_t = c$. Then $\lim_{t \rightarrow \infty} a_t/t$ exists and its value is*

$$\lim_{t \rightarrow \infty} \frac{a_t}{t} = \frac{c}{1 + b}. \quad (2.15)$$

Lemma 2.15.

- $M_1(\mathbf{B})$ exists and equals $(1 - r)/(1 + C_B)$,
- For $k \geq 2$, $M_k(\mathbf{B})$ exists and equals $M_{k-1}(\mathbf{B}) \cdot (k - 1)C_B/(1 + kC_B)$,
- $M_1(\mathbf{R})$ exists and equals $r/(1 + C_R)$, and
- For $k \geq 2$, $M_k(\mathbf{R})$ exists and equals $M_{k-1}(\mathbf{R}) \cdot (k - 1)C_R/(1 + kC_R)$,

Proof. To prove the base case, we start from Equation (2.12) for $m_{1,t}(\mathbf{B})$, use Lemma 2.13 and note that in the worst case $m_{1,t}(\mathbf{B})$ can never exceed t , and write:

$$\begin{aligned} \mathbb{E}[m_{1,t}(\mathbf{B})] &\leq \left(1 - \frac{1}{t^4}\right) \mathbb{E}[m_{1,t}(\mathbf{B})] \left(1 - \frac{C_B \pm o(1)}{t}\right) + (1 - r) + \frac{1}{t^4}t \\ &= \mathbb{E}[m_{1,t}(\mathbf{B})] \left(1 - \frac{C_B \pm o(1)}{t}\right) + (1 - r) + O\left(\frac{1}{t^3}\right). \end{aligned}$$

It only remains to apply Lemma 2.14 with $a_t = \mathbb{E}[m_{1,t}(\mathbf{B})]$, $b_t = C_B \pm o(1)$, and $c_t = (1 - r) + O(1/t^3)$.

To prove the general case, we similarly start from Equation (2.12) for $m_{k,t}(\mathbf{B})$, use Lemma 2.13, note that in the worst case $m_{k,t}(\mathbf{B})$ can never exceed t , and upper bound the expectation of $m_{k,t+1}(\mathbf{B})$ by

$$\begin{aligned} \mathbb{E}[m_{k,t+1}(\mathbf{B})] \leq & (1 - 1/t^4)\mathbb{E}[m_{k,t}(\mathbf{B})] \left(1 - \frac{(C_B \pm o(1))k}{t} \right) \\ & + (1 - 1/t^4)\mathbb{E}[m_{k-1,t}(\mathbf{B})](C_B \pm o(1))\frac{k-1}{t} + (1/t^4)t. \end{aligned}$$

It remains to apply Lemma 2.14 again, with the following three sequences: $a_t = \mathbb{E}[m_{k,t}(\mathbf{B})]$, $b_t = k(C_B \pm o(1))$, $c_t = \mathbb{E}[m_{k-1,t}(\mathbf{B})](C_B \pm o(1))\frac{k-1}{t} + O(1/t^3)$, and use induction to conclude the proof.

By applying the same techniques for the red vertices, we get the other two statements of the Lemma. \square

With that we have shown that the distribution of the red and blue nodes follow this recurrence. It is left to show that this recurrence formula leads to a power law distribution and the emergence of the glass ceiling. To show this we will need the following lemma. Using Equations 2.8 and 2.13 and lengthy algebraic conversions it is possible to show the following about C_B and C_R :

Lemma 2.16. *If $0 < r < 1/2$ and $0 < \rho < 1$ then $C_R < \frac{1}{2} < C_B$*

To show that the degree distributions of both the red and the blue vertices follow power laws we recall that a power law distribution has the following property: $M_k \propto k^{-\beta}$ for large k , where β is independent of k . If $M_k \propto k^{-\beta}$, then

$$\frac{M_k}{M_{k-1}} = \frac{k^{-\beta}}{(k-1)^{-\beta}} = \left(1 - \frac{1}{k}\right)^{\beta} = 1 - \frac{\beta}{k} + O\left(\frac{1}{k^2}\right).$$

Solving for the blue vertices, $M_k(\mathbf{B})$ and the blue exponent $\beta(\mathbf{B})$, and using Lemma 2.15, we get:

$$\frac{M_k(\mathbf{B})}{M_{k-1}(\mathbf{B})} = \frac{(k-1) \cdot C_B}{1 + k \cdot C_B} = 1 - \frac{C_B + 1}{k \cdot C_B + 1} = 1 - \frac{1 + \frac{1}{C_B}}{k} + O\left(\frac{1}{k^2}\right)$$

hence $\beta(\mathbf{B}) = 1 + 1/C_{\mathbf{B}}$. Similarly, for red vertices of degree k , $M_k(\mathbf{R})$ decays according to a power law with exponent $\beta(\mathbf{R}) = 1 + 1/C_{\mathbf{R}}$. Note that when $C_{\mathbf{R}} < \frac{1}{2} < C_{\mathbf{B}}$ we have $\beta(\mathbf{R}) > 3 > \beta(\mathbf{B})$ thus proving Theorem 2.5.

Now we have all the pieces to prove the occurrence of a glass ceiling for the biased preferential attachment process. Equipped with Theorem 2.5, Part (ii) of Theorem 2.1 follows easily. Indeed, for the *tail glass ceiling* effect, let $k(n) = n^{\frac{1}{\beta(\mathbf{R})}}$. Then

$$\mathbb{E}[\text{top}_k(\mathbf{R})] = n(\mathbf{R}) \sum_{k' \geq k} M_{k'}(\mathbf{R}),$$

$$\mathbb{E}[\text{top}_k(\mathbf{B})] = n(\mathbf{B}) \sum_{k' \geq k} M_{k'}(\mathbf{B}).$$

For $k' = n^{\frac{1}{\beta(\mathbf{R})}}$ we have $nM_{k'}(\mathbf{R}) = O(n \cdot n^{-\frac{\beta(\mathbf{R})}{\beta(\mathbf{R})}}) = O(1)$ while it holds that $nM_{k'}(\mathbf{B}) = \Omega\left(n \cdot n^{-\frac{\beta(\mathbf{B})}{\beta(\mathbf{R})}}\right) = \Omega(n^{1-\frac{\beta(\mathbf{B})}{\beta(\mathbf{R})}}) = \Omega(n^\epsilon)$ for $\epsilon > 0$. The result then follows since $n(\mathbf{R}) < n(\mathbf{B})$ and $M_{k'}(\mathbf{R}) < M_{k'}(\mathbf{B})$ for $k' > k$. Informally this means that looking at the nodes of degree at least $k(n) = n^{\frac{1}{\beta(\mathbf{R})}}$ the number of women is vanishingly small, even though there are many men found with a degree at least k .

For the *moment glass ceiling* effect we can show similarly:

$$g = \lim_{n \rightarrow \infty} \frac{\sum k^2 M_k(\mathbf{R})}{\sum k^2 M_k(\mathbf{B})} = \lim_{n \rightarrow \infty} \frac{O(n^{3-\beta(\mathbf{R})})}{\Omega(n^{3-\beta(\mathbf{B})})} = \lim_{n \rightarrow \infty} O\left(\frac{1}{n^{\epsilon'}}\right) = 0$$

for some $\epsilon' > 0$. These results conclude the proof for Theorem 2.1 Part (ii).

2.3.4 Proof of Theorem 2.2

In this section we prove that the glass ceiling does not emerge in our model if one of the three assumptions (a minority-majority partition, homophily, the ‘‘rich get richer’’ mechanism) is not met.

No Minority

No minority means that the entry rate of the red nodes is the same as the entry rate of the blue nodes, namely $r = 1/2$. We prove the following lemma:

Lemma 2.17. *If $r = 1/2$ then $C_R = C_B = 1/2$.*

Proof. Setting $r = 1/2$ in Equations 2.13 we get the following equations:

$$\begin{aligned} C_B &= \frac{1}{2} \left(\frac{\rho}{2\alpha + 2(1-\alpha)\rho} + \frac{1}{2\alpha\rho + 2(1-\alpha)} \right) \\ C_R &= \frac{1}{2} \left(\frac{\rho}{2\alpha\rho + 2(1-\alpha)} + \frac{1}{2\alpha + 2(1-\alpha)\rho} \right). \end{aligned}$$

Substituting α with the solution to Equation 2.8 with $r = 1/2$ yields the claim. \square

Having the same exponent for both of the subpopulations we get:

$$\beta(\mathbf{B}) = \beta(\mathbf{R}) = 1 + \frac{1}{1/2} = 3 \quad (2.16)$$

This leads to the same degree distribution of the two sub-populations and no glass ceiling effect emerges.

No Homophily

Assuming no homophily, namely $\rho = 1$, means, that new nodes accept their first choice independent of the color of the chosen node. Setting $\rho = 1$ in Equation 2.13 we get the following equation:

$$C_B = C_R = 1/2, \quad (2.17)$$

hence the exponent is the same for both of the subpopulations. As our model without homophily is equal to the original Barabasi-Albert model (with two different colors) it is not surprising, that this leads for both subpopulations to the same exponent found for the original model. Since the degree distribution is the same for men and for women, there is no glass ceiling effect.

No Heterophily

No heterophily, meaning new nodes attach themselves only to nodes with the same color, leads to a completely segregated network. Setting $\rho = 0$ in Equation 2.13 we get the following:

$$\begin{aligned} C_{\mathbf{B}} &= \frac{1-r}{2(1-\alpha)} \\ C_{\mathbf{R}} &= \frac{r}{2\alpha} \end{aligned}$$

Using again Equation 2.8 with $\rho = 0$, we find that $\alpha = r$. Plugging these values into equation 2.18 we get the same results as without homophily.

$$C_{\mathbf{B}} = C_{\mathbf{R}} = 1/2, \quad (2.18)$$

This leads to the same degree distribution which induces that the glass ceiling effect does not emerge.

No Preferential Attachment

Let $U(n, r, \rho)$ be a random graph model similar to $G(n, r, \rho)$ except that when a new vertex z arrives at time t it chooses its neighbor $v \in V_t$ uniformly at random (with probability $1/t$). For simplicity we start at time $t = 0$ with a single vertex (with arbitrary color). Let $\tilde{M}_k(\mathbf{R})$ and $\tilde{M}_k(\mathbf{B})$ denote, respectively, the expected number of red and blue vertices of degree k in $U(n, r, \rho)$. First we prove the following Lemma.

Lemma 2.18. *For any $0 < r < 1/2$ and any $0 \leq \rho \leq 1$, the expected number of red and blue vertices of degree k in a random graph $U(n, r, \rho)$ follows a geometric distribution. In particular:*

$$\tilde{M}_k(\mathbf{R}) = r \cdot p_r \cdot (1 - p_r)^{k-1}$$

and

$$\tilde{M}_k(\mathbf{B}) = (1-r) \cdot p_b \cdot (1 - p_b)^{k-1},$$

where $p_r = \frac{1}{1+p_r^*}$, $p_b = \frac{1}{1+p_b^*}$, $p_r^* = \frac{r}{1-(1-r)(1-q)} + \frac{q(1-r)}{1-r(1-q)}$ and $p_b^* = \frac{rq}{1-(1-r)(1-q)} + \frac{1-r}{1-r(1-q)}$.

Proof. As before, let $\tilde{m}_{k,t}(\mathbf{B})$ (resp., $\tilde{m}_{k,t}(\mathbf{R})$) denote the number of blue (resp., red) vertices of degree k at time t . Again, we will derive a recurrence for $\mathbb{E}[\tilde{m}_{k,t}(\mathbf{B})]$. Consider a red vertex v and let p_{rr} be the probability that v is selected given that the new arrived vertex is a red vertex. Then $p_{rr} = \frac{1}{t} + \frac{t-1}{t}(1-r)(1-q)p_{rr}$, so $p_{rr} = \frac{1}{1-(1-r)(1-q)+(1-r)(1-q)/t} \frac{1}{t}$. If the new arrived vertex is blue then the probability that v will be selected is $p_{br} = \frac{q}{t} + \frac{t-1}{t}r(1-q)p_{br}$, so $p_{br} = \frac{q}{1-r(1-q)+r(1-q)/t} \frac{1}{t}$. Similarly, if v is blue then the probability it is selected if the new vertex is red is $p_{rb} = \frac{q}{1-(1-r)(1-q)+(1-r)(1-q)/t} \frac{1}{t}$ and $p_{bb} = \frac{1}{1-r(1-q)+r(1-q)/t} \frac{1}{t}$ is the new vertex is blue.

Let \mathcal{F}_t be the history of the process up to time t . Thus for any $k > 1$,

$$\begin{aligned} \mathbb{E}[\tilde{m}_{k,t+1}(\mathbf{B})|\mathcal{F}_t] &= \tilde{m}_{k,t}(\mathbf{B})(1-rp_{rb} - (1-r)p_{bb}) \\ &\quad + \tilde{m}_{k-1,t}(\mathbf{B})(rp_{rb} + (1-r)p_{bb}). \end{aligned}$$

For $k = 1$ we have

$$\mathbb{E}[\tilde{m}_{1,t+1}(\mathbf{B})|\mathcal{F}_t] = \tilde{m}_{1,t}(1-rp_{rb} - (1-r)p_{bb}) + (1-r).$$

Taking the expectation on both sides we have

$$\begin{aligned} \mathbb{E}[\tilde{m}_{k,t+1}(\mathbf{B})] &= \mathbb{E}[\tilde{m}_{k,t}(\mathbf{B})](1-rp_{rb} - (1-r)p_{bb}) \\ &\quad + \mathbb{E}[\tilde{m}_{k-1,t}(\mathbf{B})](rp_{rb} + (1-r)p_{bb}). \end{aligned}$$

$$\mathbb{E}[\tilde{m}_{1,t+1}(\mathbf{B})] = \mathbb{E}[\tilde{m}_{1,t}](1-rp_{rb} - (1-r)p_{bb}) + (1-r).$$

Let

$$\tilde{M}_k(\mathbf{B}) = \lim_{t \rightarrow \infty} \frac{\mathbb{E}[\tilde{m}_{k,t}(\mathbf{B})]}{t}. \quad (2.19)$$

and

$$p_b^* = \lim_{t \rightarrow \infty} (rp_{rb} - (1-r)p_{bb})t = \frac{rq}{1-(1-r)(1-q)} + \frac{1-r}{1-r(1-q)}.$$

Then using Lemma 2.14 and setting a_t, b_t and c_t accordantly we have

- $\tilde{M}_1(\mathbf{B}) = \frac{1-r}{1+p_b^*}$
- $\tilde{M}_k(\mathbf{B}) = \tilde{M}_{k-1} \frac{p_b^*}{1+p_b^*} = \tilde{M}_{k-1} \left(1 - \frac{1}{1+p_b^*}\right) = \tilde{M}_1 \left(1 - \frac{1}{1+p_b^*}\right)^{k-1}.$

Similarly, letting

$$p_r^* = \lim_{t \rightarrow \infty} (rp_{rr} - (1-r)p_{br})t = \frac{r}{1 - (1-r)(1-q)} + \frac{q(1-r)}{1-r(1-q)}$$

we get

- $\tilde{M}_1(\mathbf{R}) = \frac{r}{1+p_r^*}$
- $\tilde{M}_k(\mathbf{R}) = \tilde{M}_{k-1} \frac{p_r^*}{1+p_r^*} = \tilde{M}_{k-1} \left(1 - \frac{1}{1+p_r^*}\right) = \tilde{M}_1 \left(1 - \frac{1}{1+p_r^*}\right)^{k-1}$

Setting $p_b = \frac{1}{1+p_b^*}$ and $p_r = \frac{1}{1+p_r^*}$ the result follows. \square

Theorem 2.6. *For any $0 < r < 1/2$ and any $0 \leq \rho \leq 1$, the random graph $U(n, r, \rho)$ does not exhibit a strong glass ceiling effect.*

Proof. Since both degree distributions follow a geometric distribution we have:

$$g = \lim_{n \rightarrow \infty} \frac{\sum k^2 \tilde{M}_k(\mathbf{R})}{\sum k^2 \tilde{M}_k(\mathbf{B})} = \frac{r(1-2p_r)/p_r^2}{(1-r)(1-2p_b)/p_b^2} = \Omega(1).$$

\square

2.3.5 Testing Competing Explanations: Unequal Qualification

Another possible competing explanation to the glass ceiling effect may be based on the conjecture that the effect occurs in areas where women have lower qualifications and skills than men. As we interpret the degree of a vertex as representing its power, unequal qualifications can be modeled by assuming that when a red (minority) vertex joins the network, it does so with a lower degree (fewer new edges) than does a blue vertex. This provides a “trivial” explanation to influence inequality, but will it cause a glass ceiling effect? We show that this is not the case: assuming no homophily, even if the minority has lower average degree, no glass ceiling effect emerges.

To formally model unequal qualifications, consider a random model similar to the *unbiased* preferential attachment model $G(n, r, \rho = 1)$ that

we denote by $G_\Delta(n, r)$, operating as follows. At each time t a new vertex w joins the graph. Its color is red with probability r and blue with probability $(1 - r)$. If w is red, then it generates one new edge according to preferential attachment as before. If w is blue then it generates Δ new edges, one at a time, according to preferential attachment.

We prove the following.

Theorem 2.7. *Let $0 < r < \frac{1}{2}$ and Δ be a constant integer. Then $G_\Delta(n, r)$ does not exhibit a tail glass ceiling effect. Formally, for every k s.t. $\lim_{n \rightarrow \infty} \text{top}_k(\mathbf{B}) = \infty$,*

$$\lim_{n \rightarrow \infty} \frac{\text{top}_k(\mathbf{R})}{\text{top}_k(\mathbf{B})} > c,$$

where $c > 0$ is a constant that depends only on r and Δ .

Proof. Instead of $G_\Delta(n, r)$, let us consider an equivalent process $G_1^\Delta(n, r)$ defined as follows. First generate $G_1(\Delta n, r)$ without coloring the vertices, according to the preferential attachment model. Then, consider the vertices of $G_1(\Delta n, r)$ in order of arrival, $v^1, v^2, \dots, v^{\Delta n}$. Generate $G_1^\Delta(n, r)$ and its vertices v_1, v_2, \dots, v_n as follows. Initially $j = i = 0$. Assume by induction that vertices v_1, v_2, \dots, v_i were already generated by processing vertices v^1, v^2, \dots, v^j . With probability r , v_{i+1} is red, in which case set the neighbors of v_{i+1} to be the vertices of G_1^Δ corresponding to the neighbors of v^{j+1} and increment j by 1. With probability $1 - r$ vertex v_{i+1} is blue, in which case set v_{i+1} to be the “merging” of $v^{j+1}, v^{j+2}, \dots, v^{j+\Delta}$ into a single blue vertex. That is, the set of neighbors of v_{i+1} is taken to be the union the sets of vertices of G_1^Δ corresponding to the neighbors of $v^{j+1}, v^{j+2}, \dots, v^{j+\Delta}$; then increment j by Δ . Once n vertices are generated in this way, we ignore all remaining vertices $v^{j'}$ that were not used, as well as their edges. This defines $G_1^\Delta(n, r)$.

For the analysis, let n^* denote the number of vertices used from $G_1(\Delta n, r)$ and note that $n \leq n^* \leq \Delta n$. Let $G_1(\Delta n, r)[n^*]$ denote the induced subgraph of $G_1(\Delta n, r)$ from vertices v^1, v^2, \dots, v^{n^*} . We prove that $G_1^\Delta(n, r)$ does not exhibit a glass ceiling effect and therefore $G_\Delta(n, r)$ doesn't either.

Let $\text{top}_k(\mathbf{R}, G)$ denote the expected number of red vertices of degree at least k in the graph G . Consider a red vertex of degree k in $G_1^\Delta(n, r)$.

The expected number of red vertices of degree at least k in $G_1^\Delta(n, r)$ is the same as in $G_1(\Delta n, r)[n^*]$.

Let $\hat{M}_k(\mathbf{R})$ and $\hat{M}_k(\mathbf{B})$ denote, respectively, the expected number of red and blue vertices of degree k in $G_1(\Delta n, r)[n^*]$. Note that $G_1(\Delta n, r)[n^*]$ follows a power law for both the red and blue vertices, and with the same β , namely, $\hat{M}_k(\mathbf{R}) \propto k^{-\beta}$ and $\hat{M}_k(\mathbf{B}) \propto k^{-\beta}$ for large k . Hence for constants c' and c'' we have:

$$\begin{aligned} \text{top}_k(\mathbf{R}, G_1^\Delta(n, r)) &= \text{top}_k(\mathbf{R}, G_1(\Delta n, r)[n^*]) \\ &= \frac{r}{r + \Delta(1-r)} \text{top}_k(G_1(\Delta n, r)[n^*]) \\ &\geq \frac{r}{r + \Delta(1-r)} \text{top}_k(G_1(\Delta n, r)[n]) \\ &= \frac{r}{r + \Delta(1-r)} \int_k^n n \cdot c' \cdot i^{-\beta} di \\ &= \frac{r}{r + \Delta(1-r)} \frac{c'' \cdot n}{\beta - 1} k^{-(\beta-1)}. \end{aligned}$$

Similarly, the number of blue vertices of degree at least k in $G_1^\Delta(n, r)$ can be upper bounded by the number of blue vertices of degree at least k/Δ in $G_1(\Delta n, r)[n^*]$, since merging Δ blue vertices of degree less than k/Δ cannot yield a blue vertex of higher degree than k . For constants c' and c'' we have:

$$\begin{aligned} \text{top}_k(\mathbf{B}, G_1^\Delta(n, r)) &\leq \text{top}_{\frac{k}{\Delta}}(\mathbf{B}, G_1(\Delta n, r)) \\ &= \frac{\Delta(1-r)}{r + \Delta(1-r)} \text{top}_{\frac{k}{\Delta}}(G_1(\Delta n, r)) \\ &= \frac{\Delta(1-r)}{r + \Delta(1-r)} \int_{\frac{k}{\Delta}}^{\Delta n} \Delta n \cdot c' \cdot i^{-\beta} di \\ &= \frac{\Delta(1-r)}{r + \Delta(1-r)} \frac{c'' \cdot \Delta n}{\beta - 1} k^{-(\beta-1)} \Delta^{\beta-1}. \end{aligned}$$

Putting it all together, we have that in $G_1^\Delta(n, r)$ there is no *tail* glass ceiling effect:

$$\frac{\text{top}(\mathbf{R})_k}{\text{top}(\mathbf{B})_k} \geq \frac{r}{(1-r)\Delta^{\beta+1}}.$$

A similar argument shows that there is also no *strong* glass ceiling effect. \square

2.3.6 Independence of tail and moment glass ceiling effects

To establish the independence of the two definitions of the glass ceiling effect, consider two sets of degree sequences, denoted A and B , where each set contains two degree sequences, for the red and blue vertices respectively. For simplicity, each degree sequence is of size n and the combined graph has $2n$ vertices. For each such set it is easy to construct a graph with the given degree sequences, for example by the random configuration model (which generates a random graph for every given degree sequence). Set A exhibits a tail glass ceiling effect but not a strong moment glass ceiling effect, whereas set B exhibits a strong glass ceiling effect but not a tail glass ceiling effect.

Set A. Let the degree sequence of the red vertices consist of $n - \sqrt{n}$ vertices of degree 1 and \sqrt{n} vertices of degree $h = \lfloor \log n \rfloor$. The degree sequence of the blue vertices consists of $n - \sqrt{n}$ vertices of degree 1 and \sqrt{n} vertices of degree $3h$.

Taking $k = 2h$, we get $\lim_{n \rightarrow \infty} \text{top}_k(\mathbf{B}) = \infty$ and

$$\lim_{n \rightarrow \infty} \frac{\text{top}_k(\mathbf{R})}{\text{top}_k(\mathbf{B})} = \frac{0}{\sqrt{n}} = 0.$$

However, the network does not exhibit a strong moment glass ceiling effect, as

$$\lim_{n \rightarrow \infty} \frac{\frac{1}{n(\mathbf{R})} \sum_{v \in \mathbf{R}} \delta(v)^2}{\frac{1}{n(\mathbf{B})} \sum_{v \in \mathbf{B}} \delta(v)^2} = \lim_{n \rightarrow \infty} \frac{n - \sqrt{n} + \sqrt{n}h}{n - \sqrt{n} + 9\sqrt{n}h} = 1 - o(1) \geq \frac{1}{2}.$$

Set B. Let the degree sequence of the red vertices consist of n vertices of degree 2 (e.g., a ring). The degree sequence of the blue vertices consists of $n - 1$ vertices of degree 1 and one vertex of degree $n - 1$ (e.g., a star graph).

Taking $k > 1$, we get $\lim_{n \rightarrow \infty} \text{top}_k(\mathbf{B}) = 1$ and the condition does not hold. If we take $k = 1$ then

$$\lim_{n \rightarrow \infty} \frac{\text{top}_k(\mathbf{R})}{\text{top}_k(\mathbf{B})} = \frac{n}{n} = 1.$$

Hence there is no tail glass ceiling effect. However, the network does exhibit a strong moment glass ceiling effect, as

$$\lim_{n \rightarrow \infty} \frac{\frac{1}{n(\mathbb{R})} \sum_{v \in \mathbb{R}} \delta(v)^2}{\frac{1}{n(\mathbb{B})} \sum_{v \in \mathbb{B}} \delta(v)^2} = \lim_{n \rightarrow \infty} \frac{\sum_1^n 2^2}{n-1 + (n-1)^2} = \lim_{n \rightarrow \infty} \frac{4n}{n(n-1)} = 0.$$

2.4 Empirical Results

To provide empirical evidence illustrating the results of our analysis in real-life, we studied a *mentor-student network* of researchers in computer science. We extracted this data from DBLP [58], a dataset recording most of the publications in computer science. A filtering process creates a list of edges connecting students to mentors. For each edge we determined the gender of the student and of the mentor and the year in which the connection was established. The resulting network spans over 30 years and has 434232 authors and 389296 edges. In the remainder of this section we describe the data collection and the assignment of gender to names followed by a temporal analysis of the minority-majority partition, the influence inequality and the glass ceiling effects.

2.4.1 Data Collection and Gender Assignment

Assigning Gender to Names

Unfortunately, the DBLP dataset (as well as the genealogy dataset) does not include direct information about the gender of the authors. In order to determine the gender of the authors, we made use of the fact that in most languages the first name also encodes the gender. Difficulties arise with unisex names, names that lose their gender information while being translated to the Latin alphabet (such as Chinese names), as well as single letter abbreviations (such as “A. Smith”). In order to match first names with their corresponding gender we built a dictionary including first names, their corresponding gender and a number between $[0, 1]$ describing the probability of the person with this first name being assigned the correct gender.

In order to build our $\{\text{name:gender}\}$ dictionary we used four different datasets.

Dataset 1: US Birth Names (85547 names)

This dataset includes all names from Social Security card applications for newly borns in the United States between 1879 and 2012¹. For each year exists a list of names and the respective count of each gender. We summed up the counts over all years, i.e., we produced an entry for the total count of each name for each gender. From this list a name was assigned to be female or male respectively if its gender's count was more than 90% of the total. Otherwise the name was assumed to fit both gender. The female score of a name would thus be $count_f / (count_f + count_m)$.

Dataset 2: US Census 1990 Data (5163 names)

This dataset has been composed by the Census Bureau in the US² and contains the names and the frequency of names for the sample male and female population respective according to the 1990 census. As an example: If, in the population of the 1990 census, 2000 people were named "Patricia" and 1900 were female, we assign it the gender "female" with probability 0.95 (score).

Dataset 3: Popular US Baby Names (4411 names)

This dataset stems from from the US Social Security Administration's statistics³ for popular baby names and contains for every year between 1960 and 2010 the 100 most popular baby names. For each year and name the average probability of usage between 1960 and 2010 was calculated. In order to assign a name its gender, we compared the male probability to the female probability and assigned the conditional probability to the gender with the higher probability.

Dataset 4: Baby Name Lists for Parents (19833 names)

The last dataset used is based on the gender information on a homepage collecting information on names made to help parents to choose a name for their baby⁴. This site notes for each name whether it is used as a female, male or unisex name. However, there is no information on the frequency of its use for each gender. The score for a female only or male only name is hence set to 1.

The datasets mentioned above were unified via the following protocol.

¹<http://www.ssa.gov/OACT/babynames/limits.html>

²http://www.census.gov/genealogy/www/data/1990surnames/names_files.html

³<http://www.ssa.gov/cgi-bin/popularnames.cgi>

⁴<http://www.behindthename.com>

First, generate a list L based on Dataset 1. Each entry of the list consists of the tuple $(name, gender, score)$. Second, process Dataset 2: for each name of Dataset 1 that is already in L , check if the gender assignment is equivalent in both data sets. If so, the list is not modified. If the gender differs, then the name is declared to be unisex. If the name is not in L , add it together with the score of Dataset 2. Repeat the same process with Datasets 3 and 4.

In total, these sources led to a collection of 96,314 distinct names, including 36,316 names with a score of more than 0.9 for males and 58,827 names with a score of more than 0.9 for females. To assign a gender to the authors in DBLP we looked up their first name in our dictionary. If the probability of the name being female or male was over 90%, then the corresponding gender was assigned to the author. We refer to the resulting dictionary as *version v0*.

Cross Checking and Validation of Influential Authors

To make sure that the very active authors are identified correctly and to prevent a case where a name is representing several authors with the same name (a known problem in DBLP), we carried out a number of heuristic cross checks and validations. In particular, for the top 1000 authors in the dataset, we ran a script that filtered out potentially problematic nodes. This resulted in excluding 452 (0.03%) nodes from the dataset. We refer to the resulting dictionary *version v1*. We believe that our version v1 dictionary is a cleaner and more accurate one, and therefore we present in this work results that are based on it. But in fact, the overall results of versions v0 and v1 are very similar. Table 2.1 summarizes the numbers of the gender assigning process.

Construction of Mentorship Graph

The empirical part of this article focuses on mentor relationships, as they are significantly influenced by homophily, as described in [40]. According to the Oxford Dictionary, a mentor is “an experienced person in a company or educational institution who trains and counsels new employees or students”. In a scientific context, this typically includes guidance on writing research articles, especially for the first publications. Thus, mentors are often co-authors of young researchers in the

Total Number of Authors in DBLP	1359616	100%
Female	172532	12.69%
Male	618830	45.52%
Names fits both	53330	3.92%
Excluded from top 1000	452	0.03%
Not found in name list	514472	37.84%

Table 2.1: The numbers of authors identified as female or male, or as having a name that fits both genders in the dictionary version v1. Also listed is the number of authors that have been excluded from the top 1000 authors. Note that this leads to a gender ratio of 21.08% females for the names with a gender assigned to them.

first few years of their career. Clearly, a mentor has to have some experience, i.e., the mentor has started publishing a few years earlier than the mentee. Apart from websites such as the *Mathematics Genealogy Project* (“<http://genealogy.math.ndsu.nodak.edu/>”), which collects data on PhD students and their advisors in the field of mathematics, we are unaware of databases on scientific mentoring⁵. Therefore, we base our empirical findings on a mentoring graph constructed from the publication data base DBLP, an on-line reference for bibliographic information on major computer science publications. It has evolved from an early small experimental web server to a popular open-data service for the computer science community. The entire DBLP dataset is freely available as a large XML file containing all bibliographic records. For each publication, this database provides the authors, the year, and the journal or conference, among other data.

For each author in DBLP, we looked at the set of co-authors to find potential mentors. More precisely, we considered all people that co-authored an article in the first four years of a young researcher (we only considered papers with up to 20 authors) . In addition we relied on the assumption that the mentor has significantly more experience than the mentee. We looked at the years when the researchers wrote their first article. A person was only considered as a potential mentor of a mentee if the difference

⁵we empirically studied the Mathematics Genealogy Project, with similar results, but do not report it here since the network size is much smaller.

in the number of years between the dates of their first articles exceeded four. Like this we computed a set of eligible mentoring candidates for each author in DBLP. Among these candidates, we selected the one with the highest number of early papers written (if there are several authors competing for this position, we picked one at random).

The DBLP snapshot from December 23, 2013 contains 8,867,408 articles with two or more authors, written by a total of 1,282,790 people. Among them 871,839 have a set of at least one mentor candidate, i.e., 68.01 % of the authors in DBLP can be assigned a mentor with this method.

When using the same procedure but requiring an experience difference of at least 5 or 6 years, the percentage of authors than can be assigned an author decreases to 65% and 62% respectively. The changes in our observations however is negligible. Table 2.2 gives general statistics on the mentor graph. Note that any author with degree 2 or above is a mentor.

2.4.2 Temporal Analysis

As may be expected based on previously reported studies, our mentor-student network exhibits a minority-majority partition (namely, a low proportion of up to 21% females), homophily, power law distribution and a glass ceiling effect.

Figure 2.5(a) reveals that over time, the fraction of females in the network ($n(\mathbf{R})/n$, the shaded red area) has increased, but it is still below 21%. Also the average degree for females vertices is lower (1.48 vs 1.87). Figure 2.5(b) presents an indication for homophily in the mentoring selection process. This is done by the *homophily test* of [20], which compares the expected number of “mixed” (female-male) edges to the observed one (see also Section 2.2.2).

Figure 2.6 presents indications for the glass ceiling effect. Figure 2.6(a) shows that the fraction of females among the vertices of degree k or higher, namely, $\text{top}_k(\mathbf{R})/\text{top}_k(\mathbf{B})$, decreases continuously as k increases. The first major decrease occurs when moving from the group of “students” (i.e., degree 1 vertices) to the group of researchers of degree 2 or higher: the fraction of females drops from $\text{top}_1(\mathbf{R}) \approx 21\%$ to $\text{top}_2(\mathbf{R}) < 15\%$. It is

Parameter	Value	%
Number of Females	90035	20.73
Number of Males	344197	79.27
Sum of Females and Males	434232	100.00
Mixed Edges	101607	26.09
Female-Female edges	16074	4.12
Male-Male Edges	271615	69.77
Total number of edges	389296	100.00
Sum of Female degrees	133755	17.18
Sum of Male degrees	644837	82.82
Sum of edges	778592	100.00
Number of Female Mentors	10819	14.34
Number of Male Mentors	64638	85.66
Sum of Female and Male Mentors	75457	100.00
Females Avg. Degree	1.48	
Males Avg. Degree	1.87	
Avg Degree	1.79	
Female Mentors Avg. Degree	4.60	
Male Mentors Avg. Degree	5.25	
Mentors Avg. Degree	5.16	

Table 2.2: DBLP mentor graph statistics




Figure 2.5: Female rate and *normalized power* in the computer science mentor graph. (a) The rate (i.e., percentage in the population) of females over time, compared with their normalized power, defined as $d(\mathbf{R})/(2m)$. Males have more power than expected by their rate, while females have less power than expected by their rate. (b) Evidence for homophily: a comparison of the observed number of “mixed” edges to the expected value assuming there is no homophily. We consider two cases: (i) the expected number of mixed edges ignoring the difference between the male and female average degree (expected: 127963.09 std: 293.08) and (ii) the expected number of mixed edges while considering the different degree sequences for males and females (expected: 110777.11 std: 281.52). In both cases the observed value (101607 edges) significantly deviates from the expectation (the error bars indicate the expected value ± 10 times the standard deviation) with extremely low p-values.

important to note that the data indicates that even at the high end of the graph, a few female researchers with very high degrees are still present; however, our definitions for the glass ceiling ignore this extremal effect, which is caused by a few individuals, and concentrate on the averages over large samples. Indeed, when the sample size is large enough, the fraction of the female researchers decreases. Figure 2.6(b) shows a strong indication that the degree distribution of the vertices (females, males and combined) follows a power law. This in turn is associated with a preferential attachment mechanism that is known to result in a power law degree distribution. Note that the power law exponent β for the graph of the female researchers is $\beta = 2.91$ (in the best fit), which is higher than the corresponding exponent in the graph for the male researchers, $\beta = 2.58$. Our analysis (presented in 2.3.2 and 2.3.3) establishes that if




Figure 2.6: Glass ceiling effect in mentor graph: (a) percentage of females in the mentor population of degree at least k . Female start with 21% in the population and drop to below 15% when considering degree at least 2 (faculty members). It continues to decrease (ignoring small samples at the end, see text). Vertex size and darker color represent larger sample space. (b) The power law-like degree distribution for both females and males. The exponent β for females is higher than for males, demonstrating the glass ceiling effect.

the degree distribution of both sub-populations follow a power law and the exponent for the minority sub-population is higher than that of the majority sub-population, then a strong moment glass ceiling effect will appear.

2.5 Conclusion

One obvious limitation of our model is that it is somewhat simplistic and captures only one possible mechanism for generating a glass ceiling effect. It ignores many important aspects of life (such as attraction, unequal distributed family responsibilities, and different expectations, to name a few) and alternative (co-existing) mechanisms that contribute to the effect. For instance, our model cannot be used to explain the occurrence of a glass ceiling effect in contexts where pairwise individual interactions play a less dominant role than in academia. To account for the glass ceiling effect in such contexts as well as others, one may consider alternative explanations. In particular, a common possible explanation is the

“leaky pipeline” phenomenon, namely, the phenomenon that women tend to quit or slow down their careers in order to invest more time in their families. This phenomenon can be modeled mathematically in several different ways. One such way is by introducing vertex departures in addition to vertex arrivals, with a bias in the form of increased departure rate of the minority group. But in fact, such a dynamic “leaky pipeline” model allows several reasonable sub-models that will not generate a glass ceiling effect, as well as some other sub-models that do. Moreover, the cause and effect relationship between the glass ceiling and the leaky pipeline are not necessarily one-directional; while the glass ceiling effect may indeed be the outcome of the “leaky pipeline” phenomenon in certain settings, there are other settings where it may be its (partial) cause. An interesting direction for future work would be to describe a more complete model, most likely combining a number of different mechanisms contributing jointly to the glass ceiling effect. In any case, we find it remarkable that the simple mathematical mechanism presented here (based on homophily) is sufficient to explain (at least parts of) the glass ceiling effect, despite the fact that it does not utilize the “leaky pipeline”.

Our findings may suggest ways to deal with the glass ceiling phenomenon. By better understanding the roots of the glass ceiling effect, one can address each of the elements and attempt to mitigate them or deal with those elements that are easier to manage. Our research indicates that for certain mechanisms involved in the formation of a glass ceiling, removing one element may eliminate the glass ceiling effect. Hence, while it might be difficult to modify the human tendencies of homophily and preferential attachment one could attempt to balance the proportions of minorities within the population or impose a proportional representation of successful women at the top level. Both of these options may be classified as variants of affirmative action, but the latter, even if more common, seems to avoid the roots of the problem. In particular, a more equally represented society could be created by encouraging minorities to enter the system, as our findings indicate that increasing the ratio of minorities at the entry stage may mitigate the glass ceiling effect at least partially. This conclusion is in line with a common view [74, 85], which states that fixing the “leaky pipeline” is key for a more equal gender distribution in science. By determining and examining the causes of the glass ceiling

effect, we can work on alleviating the glass ceiling effect, resulting in a richer and more diverse community.

3

Influence in a Social Network

In the last chapter we took a closer look of what effects can be observed in social networks which are built upon the inherent preferences of people. In this chapter we shed light on influence dynamics inside already existing networks. We are interested how opinions propagate through society if the opinion of people is constantly influenced by the opinion of their friends. In this chapter we abstract people as nodes in a network and their influences onto each other as edges. We call networks, in which entities of the network are continuously influenced by the states of their respective neighbors, influence networks (IN). One can find many such networks in different areas such as in social networks, belief propagation, spring embedders, cellular automata, distributed message passing algorithms, traffic networks, the brain, biological cell systems, and ant colonies, just to name a few. All of these examples of *influence networks* (INs) are known to be difficult to analyze. Some of the applications mentioned are notorious to

have long-standing open problems regarding convergence.

In this chapter we deal with a generic version of such networks: The network is given by an arbitrary graph $G = (V, E)$, and nodes of the graph switch their state to the state of the majority of their respective neighbors. We are interested in the stability of such INs with each node having a binary state. We investigate synchronous state switching as well as iterative state switching. Specifically, we would like to determine whether an IN converges to a stable situation or not. We are interested in how to specify such a stable setting, and in the amount of time needed to reach such a stable situation. We study several models how the nodes take turns, synchronous, asynchronous, adversarial, benevolent and several graph classes, unweighted, weighted, undirected and directed graphs.

The most surprising result is for synchronous INs: Each node is assigned an initial state from the set $\{R, B\}$, and in every round, all nodes switch their state to the state of the majority of their neighbors simultaneously. This specific problem is commonly referred to as “Democrats and Republicans”, see e.g. Peter Winkler’s CACM column [91]. It is well known that this problem stabilizes in a peculiar way, namely that each node eventually is in the same state every second round, which was proven by Goles and Olivos in [35]. This result can be shown by using a potential bound argument, i.e., until stabilization, in each round at least one more edge becomes “more stable”. This directly gives a $\mathcal{O}(n^2)$ upper bound for the convergence time. On the other hand, using a slightly adapted linked list topology, one can see that convergence takes at least $\Omega(n)$ rounds. But what is the correct bound for this classic problem? Most people that worked on this problem seem to believe that the linear lower bound should be tight, at least asymptotically. Surprisingly, in the course of our research, we discovered that this is not true. In this chapter we show that the upper bound is in fact tight up to a polylogarithmic factor. Our new lower bound is based on a novel graph family, which has interesting properties by itself. We hope that our new graph family might be instrumental to research concerning other types of INs, and may prove useful in obtaining a deeper understanding of some of the applications mentioned above.

We extend this model with different model variations. In particular, we look at asynchronous networks where nodes update their states se-

quentially. We show that in such a sequential setting, convergence may take $\Theta(n^2)$ time if given an adversarial sequence of steps, and $\Theta(n)$ if given a benevolent sequence of steps.

We also investigate the convergence of different graph classes such as weighted and directed graphs. We will show that in the case of weighted graphs the convergence time can grow exponentially in the number of nodes and allowing directed edges finally changes the number of possible stable states dramatically.

3.1 Related Work

Influence networks have become a central field of study in many sciences. In biology, to give three examples from different areas, [77] study networks in the context of brain science, [2] study cellular systems and their relation to distributed algorithms, and [1] study networks in the context of ant colonies. In optimization theory, belief propagation [8, 71] has become a popular tool to analyze large systems, such as Bayesian networks and Markov random fields. Nodes are continuously being influenced by their neighbors; repeated simulation (hopefully) quickly converges to the correct solution. Belief propagation is commonly used in artificial intelligence and information theory and has demonstrated empirical success in numerous applications such as coding theory. A prominent example in this context are the algorithms that classify the importance of web pages [12, 48]. In physics and mechanical engineering, force-based mechanical systems have been studied. A typical model is a graph with springs between pairs of nodes. The entire graph is then simulated, as if it was a physical system, i.e. forces are applied to the nodes, pulling them closer together or pushing them further apart. This process is repeated iteratively until the system (hopefully) comes to a stable equilibrium, [32, 42, 44, 49]. Influence networks are also used in traffic simulation, where nodes (cars) change their position and speed according to their neighboring nodes [68]. Traffic networks often use cellular automata as a basic model. A cellular automaton [69, 92] is a discrete model studied in many fields, such as computability, complexity, mathematics, physics, and theoretical biology. It consists of a regular grid of cells, each in one of a finite number of states, for instance 0 and 1. Each cell changes its state

according to the states of its neighbors. In the popular game of life [34], cells can be either dead or alive, and change their states according to the number of alive neighbors.

Our synchronous model is related to cellular automata, on a general graph; however, nodes change their opinion according to the majority of their neighbors. As majority functions play a central role in neural networks and biological applications this model was already studied during the 1980s. Goles and Olivos [35] have shown that a synchronous binary influence network with a generalized threshold function always leads to a fixed point or to a cycle of length 2. This means that after a certain amount of synchronous rounds, each participant has either a fixed opinion or changes its mind in every round. Poljak and Sura [75] extended this result to a finite number of opinions. In [36], Goles and Tchuente show that an iterative behavior of threshold functions always leads to a fixed point. Sauerwald and Sudholt [80] study the evolution of cuts in the binary influence network model. In particular, they investigate how cuts evolve if unsatisfied nodes flip sides probabilistically. To some degree, one may argue that we look at the deterministic case of that problem instead.

The exact same model as presented here, is used to study a *dynamic monopoly* [29, 62, 72], abbreviated dynamo. A dynamo is an initial set of vertices in an influence network, all with the same opinion, such that after a finite number of steps all nodes in the network share that opinion. The minimum size of a dynamo was studied in [73], where it is shown that $\Omega(\sqrt{n})$ nodes are needed for a monotone dynamo (assuming no node ever changes back its state) and for 2-round dynamos (where the network stabilizes after exactly 2 rounds). Berger [7] extends these results by proving that a constant number of nodes may suffice to convert a network of size at least n for arbitrary n . In contrast, in the current work we ignore the final opinions of the network, and focus on stabilization time.

In sociology, understanding social influence (e.g. conformity, socialization, peer pressure, obedience, leadership, persuasion, sales, and marketing) has always been a cornerstone of research, e.g. [45]. More recently, with the proliferation of online social networks such as Facebook, the area has become en vogue, e.g. [4, 65]. Leskovec et al. [56] for instance verify the balance theory of Heider [37] regarding conformity of opinions; they study how positive (and negative) influence links affect the structure of

the network. Closest to our paper is the research dealing with influence, for instance in the form of sales and marketing. For example, [55] investigate a large person-to-person recommendation network, consisting of four million people who made sixteen million recommendations on half a million products, and then analyze cascades in this data set. Cascades can also be studied in a purely theoretical model, based on random graphs with a simple threshold model which is close to our majority function [88]. Rumor spreading has also been studied algorithmically, using the random phone call model, [19, 43, 79]. Using real data from various sources, [5] show that networks generally have a core of influential (elite) users. In contrast to our model, nodes cannot change their state back and forth, once infected, a node will stay infected. Plenty of work was done focusing on the prediction of influential nodes. One wants to find subset of influential nodes for viral marketing, e.g. [14, 46]. In contrast, [50] studies the case of competitors, which is closer to our model since nodes can have different opinions. However, also in [50] nodes only change their opinion once. However, in all these social networks the underlying graph is fixed and the dynamics of the stabilization process takes place on the changing states of the nodes only. An interesting variant changes the state of the edges instead. A good example for this is matching. A matching is (hopefully) converging to a stable state, based on the preferences of the nodes, e.g. [30, 33, 47]. Hofer takes these edge dynamics one step further, as not only the state of the edge changes, but the edge itself [39].

3.2 Model

We model an *influence network* (IN) as a graph $G = (V, E, \omega, \mu_0)$. The set of nodes V is connected by an arbitrary set of edges E . Each edge is assigned a weight $\omega(e) \in \mathbb{N}$. In Sections 3.3, 3.4, and 3.5 where we talk about unweighted graphs we assume each edge having a weight of one and leave out the weight function in the description. We refer to an edge between nodes u and v with weight ω as (u, v, ω) . Usually we talk about undirected edges, except in Section 3.7 about asymmetric graphs, where we consider directed edges. In said section (u, v, ω) means there is an edge from node u towards node v with weight ω . The weight of a

graph G is defined as $\omega(G) = \sum_{e \in E} \omega(e)$. Each node has an initial opinion (or state) $\mu_0(v) \in \{R(ed), B(lue)\}$ (graphically interpreted as the Red and Blue colors respectively). The opinions of the nodes at every time t are also represented in the same way by a function $\mu_t : V \mapsto \{R, B\}$.

We define the “red neighborhood” of node v by $\Gamma_{R,t}(v) = \{u \in \Gamma(v) \mid \mu_t(u) = R\}$ at time t and similarly for the “blue neighborhood” of v as $\Gamma_{B,t}(v)$. A node changes its opinion on time step t if a weighted majority of its neighbors has a different opinion. One can consider different actions in case of a tie. We chose the nodes own opinion as a tie breaker because of two reasons. First it seems to be a natural choice and secondly one can build an equivalent graph with the same behavior in asymptotic running time by cloning the graph and connecting each node with its clone and the neighbors of its clones. More formally the state of a node at time $t+1$ is defined as

$$\mu_{t+1}(v) = \begin{cases} R, & \text{if } \sum_{u \in \Gamma_{R,t}(v)} \omega(u, v) > \sum_{u \in \Gamma_{B,t}(v)} \omega(u, v) \\ B, & \text{if } \sum_{u \in \Gamma_{B,t}(v)} \omega(u, v) > \sum_{u \in \Gamma_{R,t}(v)} \omega(u, v) \\ \mu_t(v), & \text{otherwise .} \end{cases}$$

We study both synchronous and asynchronous INs. A synchronous IN evolves over a series of rounds. In each round the nodes simultaneously update their opinion to the weighted majority of their neighbors according to the above rule.

As will be explained in Section 3.3, the only interesting asynchronous model is the sequential model. In this model, we call the change of opinion of one node a *step*. The opinion of node v after t steps is defined as $\mu_t(v)$. In general, more than one node may be ready to take a step. Depending on whether we want convergence to be fast or slow, we may choose different nodes to take the next step. If we aim for fast convergence, we call this the *benevolent sequential model*. Slow conversion on the other hand we call the *adversarial sequential model*.

As INs are deterministic, they necessarily enter a cyclic pattern after a certain number of rounds. We call an IN *stable* with a cycle of length q if each node changes its opinion in a cyclic pattern with cycle length k




Figure 3.1: An IN with a certain starting assignment converges after certain number of rounds/steps (c, c', c'') to a stable state with a cycle length of q

for some $k \leq q$. In other words, a state can be stable even though some nodes still change their opinion, see Figure 3.1 for an illustration.

Definition 3.1. An IN $G = (V, E, \omega, \mu_0)$ is *stable* at time t with cycle length q , if for all vertices $v \in V : \mu_{t+q}(v) = \mu_t(v)$. A *fixed state* of an IN G is a stable state with cycle length 1. The *stabilization time* c of an IN G is the smallest t for which G is stable.

Note that since INs are deterministic an IN which has reached a stable state will stay stable.

We investigate the stability, the convergence time c and the periodicity q of INs in the described models. An IN and its convergence process can be represented by a finite state machine as illustrated in Figure 3.1. Clearly, the convergence process depends not only on the graph structure, but also on the initial opinions of the nodes. We investigated graphs and initial opinions that maximize convergence time. In the benevolent sequential model in particular, we investigate graphs and sets of initial opinions leading to the worst possible convergence time, given the respectively best sequence of steps.

3.3 Sequential Influence Network

For simplicity we start to analyze a unweighted graph where the nodes change opinions asynchronously. In an asynchronous setting, nodes can take steps independently of each other, i.e. subsets of nodes may reassess and change their opinion concurrently. Unfortunately, in such a setting, convergence time is not well defined. To see this, consider a star-graph where the center has a different initial opinion than the leaves. An adversary may arbitrarily often chooses the set of all nodes to reassess their opinion. After r such rounds the adversary chooses only the center node. Now this IN stabilizes, after r rounds for an arbitrary $r \rightarrow \infty$. In other words, asynchrony in its most general form is not well defined in this setting, and we restrict ourselves to sequential steps only, whereas a step is a single node changing its opinion. The sequence of steps is chosen by an adversary which tries to maximize the convergence time. Note that the convergence upper bound presented in Lemma 3.2 implies immediately that the IN stabilizes in a fixed state.

Lemma 3.2. *A sequential IN reaches a fixed state after at most $\mathcal{O}(n^2)$ steps.*

Proof. Divide the nodes into the following two sets according to their current opinion: $S_R = \{v \mid \mu(v) = R\}$ and $S_B = \{v \mid \mu(v) = B\}$. If a node changes its opinion, it has more neighbors in the opposite set than in its current set. Therefore the number of edges $X = \{\{u, v\} \mid u \in S_R, v \in S_B\}$ between nodes in set S_R and set S_B is strictly decreasing. Each change of opinion reduces the number of edges of X by at least one. Therefore the number of steps is bounded by the number of edges in X . In a graph G with n nodes $|X|$ is at most $n^2/4$, therefore at most $\mathcal{O}(n^2)$ steps can take place until the IN reaches a fixed state. \square

It is more challenging to show that this simple upper bound is tight. We show a graph G and a sequence of steps in which way an adversary can provoke $\Omega(n^2)$ convergence time.

Lemma 3.3. *There is a family of INs with n vertices such that a fixed state is reached after $\Omega(n^2)$ steps.*

```

S ← ()
for i = 0 to n/3 do
  S = reverse(S);
  S ← (i, S);
  for all x ∈ S do
    take step x;
  end for
end for
    
```

Algorithm 3.1: Adversarial Sequence

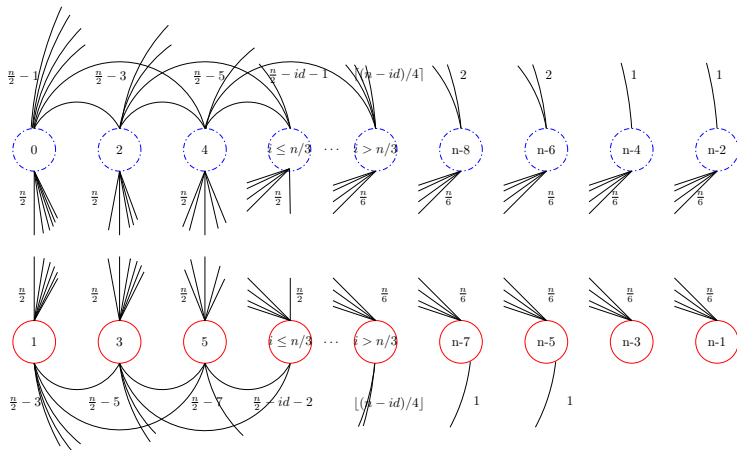


Figure 3.2: In this graph an adversary can provoke $\Omega(n^2)$ changes of opinion.

Proof. Consider the following graph G with n nodes. The nodes are numbered from 0 to $n-1$, whereas nodes with an even id are initially assigned opinion B and nodes with an odd id are assigned opinion R . See also Figure 3.2. All even nodes with $id \leq n/3$ are connected to all odd nodes. All odd nodes with $id \leq n/3$ are connected to all even nodes respectively. In addition an even node with $id \leq n/3$ is connected to nodes $\{0, 2, 4, \dots, n - 2 \cdot id - 2\}$, respectively an odd node with $id \leq n/3$ is connected to nodes $\{1, 3, 5, \dots, n - 2 \cdot id - 3\}$. For example, node 0 is a neighbor of all nodes, whereas node 1 is neighbor of all nodes except the nodes $n-1$ and $n-3$. Note that each node i with $i \leq n/3$ is connected to all other nodes with $id \leq n/3$. For each node v the change potential indicates of how much a node would like to change its opinion. It is the number of neighbors with a different opinion minus the number of neighbors who share the opinion. Formally $P(v)$ is defined as:

$$P(v) = |\Gamma_{\neq \mu(v), t}(v)| - |\Gamma_{\mu(v), t}(v)|$$

Put differently, if the change potential of a node is larger than 0, and it is requested to reassess its opinion, it takes a step. A large change potential of a node v , means that many neighbors of v have the opposite opinion from v . If a neighbor of v with the same opinion takes a step, v 's change potential $P(v)$ is increased by 2. On the other hand, if a neighbor changes from the opposite opinion to the same opinion as node v , $P(v)$ is decreased by 2. If v itself changes its opinion, its change potential turns from p to $-p$. The change potential of v is basically the number of edges by which the total number of edges between set S_B and set S_R is reduced if v changes its opinion. As the total amount of steps is bounded by the number of edges between set S_B and S_R , a node v with $P(v) = p$ reduces the remaining number of possible changes by p if it takes a step. E.g. in the previously constructed graph G , the first nodes have the following change potential: $P(0) = 1, P(1) = 3, P(2) = 3, P(3) = 5$. Generally, node i has a change potential $P(i) = n/2 - (n/2 - i - 1) = i + 1$ if i is even respectively $P(i) = i + 2$ if i is odd. In order to provoke as many steps as possible, the adversary selects the nodes which have to reassess their opinion according to the following rule: He chooses the node with the smallest id for which $P(v) = 1$. Therefore each step reduces the remaining number of possible steps by 1. G is constructed in such a way, that a step

from a node triggers a cascade of steps from nodes which have already changed their opinion whereas each change reduces the overall potential by 1.

The adversary chooses the nodes in phases according to algorithm 3.1. Phase i starts with the selection of node i followed by the selections of all nodes with $id < i$, where the adversary chooses the nodes in the reverse order than it did in round $i - 1$. Phase 0 consists of node 0 changing its opinion, in phase 1 node 1 and then node 0 make steps, and in phase 2 the nodes change in the sequence 2, 0, 1. As a node v can only change its opinion if $P(v) > 0$, we need to show that this is the case for each node v which is selected by the adversary. It is sufficient to show that each node which is selected has a change potential of 1.

We postulate:

- (i) At the beginning of phase i , the following holds: $P(i) = 1$ and $\forall v < i : \mu(v) = \mu(i)$.
- (ii) Each node the adversary selects has change potential 1 and each node with $id \leq i$ is selected eventually in phase i .
- (iii) At the end of phase i , all nodes with $id \leq i$ have opinion R if i is even and opinion B if i is odd.

We prove (i), (ii) and (iii) by induction. Initially, part (i) holds, as no node with $id < 0$ exists and as node 0 is connected to $n/2$ nodes with opinion R and to $n/2 - 1$ nodes with opinion B and therefore has change potential 1. In phase 0 only node 0 is selected, therefore part (ii) holds as well. Node 0 changed its opinion and has therefore at the end of phase 0 opinion R , therefore part (iii) holds as well.

Now the induction step: To simplify the proof of part (i) we consider odd and even phases separately. Consider an odd phase i . At the start of phase i , no node with $id \geq i$ has changed its opinion yet. Therefore node i still has its initial opinion $\mu(i) = R$. According to (iii), each node with $id \leq i - 1$ has at the end of phase $i - 1$ opinion $R = \mu(i)$. So $(i + 1)/2$ neighbors of i have compared to the initial state, changed their opinion from B to R . If a neighbor u of a node v with a different opinion than v changes it's opinion, v 's change potential is decreased by 2. Therefore node i 's initial change potential $P_{t_0}(i) = n/2 - (n/2 - i - 2) = i + 2$ is

decreased by $2 \cdot (i+1)/2 = i+1$ and is therefore $P(i) = i+2 - (i+1) = 1$ at the beginning of phase i . Therefore (i) holds before an odd phase.

Now consider an even phase i . At its start, all nodes with $id \geq i$ still have their initial opinion. Therefore node i has opinion $\mu(i) = B$. According to (iii) each node with $id \leq i-1$ has at the end of phase $i-1$ opinion $B = \mu(i)$. As node i 's initial change potential was $P_{t_0}(i) = n/2 - (n/2 - i - 1) = i+1$ and $i/2$ neighbors of i changed from opinion R to opinion B compared to the initial state, i 's new change potential is calculated as $P(i) = i+1 - 2 \cdot i/2 = 1$. Therefore (i) holds before an even phase, hence (i) holds.

To prove part (ii) let v be the last node which was selected in phase $i-1$. As v was selected, it had according to (ii) a change potential of 1. If a node changes its opinion, its change potential gets inversed. Therefore node v had at the beginning of phase i a change potential of -1 . In addition, node v is by construction a neighbor of node i and has according to (i) at the start of phase i the same opinion as node i . As node i changes its opinion, node v 's change potential is increased by 2. Therefore v 's new change potential is again $-1+2 = 1$, when it is selected by the adversary. The same argument holds for the second last selected node u . After it was selected in phase $i-1$ its change potential was -1 . Then v has changed its opinion which led to $P(u) = -3$. As node i and node v changed their opinions in phase i , $P(u)$ was again 1. Hence if the adversary selects the nodes in the inverse sequence as in phase $i-1$, each selected node has a change potential of 1 and is selected eventually. Therefore (ii) holds.

As node i and all nodes with $id \leq i-1$ had at the beginning of phase i the opinion $\mu(i)$ according to (iii) and all nodes have changed their opinion in phase i according to (ii), all nodes with $id \leq i$ must have the opposite opinion at the end of phase i , namely R if i is even or B otherwise. Therefore (iii) holds as well.

We now have proven that in phase i , i nodes change their opinion. As the adversary starts $n/3$ phases, the total number of steps is $1/2 \cdot n/3 \cdot (n/3 - 1) \in \Omega(n^2)$. \square

Directly from Lemma 3.2 and Lemma 3.3, we get the following theorem.

Theorem 3.1. *A worst case sequential IN reaches a fixed state after $\Theta(n^2)$ steps.*

We have seen, that with an adapted graph and an adversary using Algorithm 3.1 an IN takes up to $\Theta(n^2)$ steps until it stabilizes. But how bad can it get, if the process is benevolent instead?

Theorem 3.2. *An IN with a benevolent sequential process reaches a fixed state after $\Theta(n)$ steps.*

Proof. A benevolent process needs $\Omega(n)$ steps to reach a stable state. This can be seen by considering the complete graph K_n with initially $\lfloor n/2 \rfloor - 1$ red nodes and $\lceil n/2 \rceil + 1$ blue nodes. Independently of the chosen sequence this IN needs exactly $\lfloor n/2 \rfloor - 1$ steps to stabilize because the only achievable stable state is all nodes being blue. To prove that the number of steps is bounded by $\mathcal{O}(n)$ we define the following two sets: The set of all red nodes which want to change: $C_R = \{v \mid \mu(v) = R \wedge P(v) > 0\}$ and the set of all blue nodes which want to change: $C_B = \{v \mid \mu(v) = B \wedge P(v) > 0\}$. A benevolent process chooses nodes in two phases. In the first phase it chooses nodes from C_B until the set is empty. During this phase, it may happen that additional nodes join C_B (e.g. a leaf of a node $v \in C_B$, after v made a step). However, no node which left C_B will rejoin, as those nodes turned red and can not turn blue again in this phase. In the second phase, the benevolent process chooses nodes from C_R until this set is empty. The set C_B will stay empty during the second phase since nodes turning blue can only reinforce blue nodes in their opinion. Both phases take at most n steps, therefore proving our upper bound. \square

3.4 Synchronous Influence Network

In this section we consider the synchronous model, sometimes also known as “Democrats and Republicans”. In particular we investigate an unweighted graph, i.e., all the edges of the IN have weight exactly 1, and the nodes reassess their opinions simultaneously in rounds. We hereafter refer to such an IN as an *unweighted influence network*, or UIN in short. A way to think of this is, that people express their opinion (voting for the democrats party or the republican) every day by wearing either a blue or

a red hat. At the end of each day they call their friends and ask which hat their friends are wearing today. For the next day, they then choose the color of the hat according to the opinion they heard the most. We call these networks synchronous IN's. Opposed to the sequential IN, a synchronous IN may stabilize in a state where some nodes change their opinion in every round. For example, consider the graph K_2 (two nodes, connected by an edge) where the first node has opinion B and the second node has opinion R . After one round, both vertices have changed their state, which leads to a symmetric situation. This IN remains in this stable state forever with a period q of length 2. As has already been shown in [35,91], a synchronous IN always reaches a stable state with a periodicity of at most 2 after $\mathcal{O}(n^2)$ rounds.

Theorem 3.3 ([35]). *The cycle length of the stable state of a synchronous unweighted influence network with symmetric weights is at most 2.*

Theorem 3.4 ([91]). *An n -node unweighted synchronous influence network stabilizes in $\mathcal{O}(n^2)$ time steps.*

The proof of the upper bound in Theorem 3.4 uses a bound argument on the edges. Each edge (v, u) is substituted by two directed edges $\langle v, u \rangle$ and $\langle u, v \rangle$, with the same weight, referred to as the outgoing and incoming edges of v , respectively. One can think of these edges as representing “advice” given between neighbors. The outgoing edge from node v to node u can be seen as the opinion that node v proposes to its neighbor u and the incoming edge can be seen as the opinion that node u proposes to v . In each time step t , each of these directed edges is declared to either “succeed” or “fail”. The outgoing edge $\langle v, u \rangle$ succeeds on time step t if the neighbor u accepts the opinion proposed by v during the round leading from time step t to time step $t + 1$, namely, $\mu_{t+1}(u) = \mu_t(v)$, and fails otherwise. The analysis is based on the initial observation that a UIN starts with a certain number of failing edges $f(0)$ at time 0, which is naturally bounded by $f(0) \leq 2|E|$. It is shown that as long as the UIN has not stabilized, the number of failing edges $f(t)$ decreases in every round by at least one. Using the same arguments, the upper bound for a UIN can be restated as $2|E|$.

Theorem 3.5 ([91]). *An n -node unweighted influence network with edge set E stabilizes in at most $2|E|$ time steps.*




Figure 3.3: A transistor $T(4)$. The dotted lines indicate how the transistor will be connected.

In this section, we prove this bound to be almost tight.

Theorem 3.6. *There is a family of synchronous INs with convergence time $\Omega\left(\frac{n^2}{(\log \log n)^2}\right)$.*

As the the technical proof of Theorem 3.6 is quite involved we do start with a simpler IN with convergence time $\Omega(n^{3/2})$.

The basic idea is to construct a mechanism which forces vertices on a simple path to change their opinion one after the other. Every time the complete path has changed, the mechanism should force the vertices of the path to change their opinions back again in the same order. To create this mechanism, we introduce an auxiliary structure called transistor, which is depicted in Figure 3.3.

Definition 3.4. A *transistor* of size s , denoted as $T(s)$, is a graph consisting of s collector vertices $\mathcal{C} = \{v_i^C \mid 0 \leq i \leq s-1\}$, s emitter vertices $\mathcal{E} = \{v_i^E \mid 0 \leq i \leq s-1\}$ and three base vertex $\mathcal{B}^1 = v_1^B$, $\mathcal{B}^2 = v_2^B$ and $\mathcal{B}^3 = v_3^B$. All edges between collector and emitter vertices, all edges between any two base vertices, and all edges between collector vertices

and the third base vertex exist. Formally:

$$\begin{aligned}
 T(s) &= (V, E) \\
 V &= \mathcal{C} \cup \mathcal{E} \cup \{\mathcal{B}^1, \mathcal{B}^2, \mathcal{B}^3\} \\
 E &= \{\{u, v\} \mid u \in \mathcal{C}, v \in \mathcal{E}\} \cup \{\{u, \mathcal{B}^3\} \mid u \in \mathcal{C}\} \cup \\
 &\quad \{\{u, v\} \mid u, v \in \{\mathcal{B}^1, \mathcal{B}^2, \mathcal{B}^3\}, u \neq v\}
 \end{aligned}$$

All vertices in the same transistor are initialized with the same opinion $X \in \{R = 1, B = -1\}$. All collector edges that are not connected to emitter vertices, the edges from \mathcal{B}^3 that are not connected to collector or other base vertices, as well as one edge for \mathcal{B}^1 and \mathcal{B}^2 , represented by dotted edges pointing to the top in Figure 3.3, are connected to vertices with a constant opinion $-X$. We call all edges connecting the collector nodes with the constant opinion *collector edges*. We also call the edges from emitter vertices that are not connected to collector vertices *emitter edges*. They are represented as dotted edges pointing to the bottom in Figure 3.3 and may connect the emitter vertices to any vertex. We denominate the two edges from base vertex \mathcal{B}^1 that neither connect to the constant opinion nor to any other base vertex as base edges (see dotted edges pointing to the left in Figure 3.3). As soon as both base edges advertise opinion $-X$, the transistor will completely flip to opinion $-X$ in 4 rounds regardless of what is advertised over the emitter edges, i.e., the following sets of vertices will all change their opinion to $-X$ in the given order: $\{\mathcal{B}^1\}$, $\{\mathcal{B}^2, \mathcal{B}^3\}$, \mathcal{C} , \mathcal{E} .

We now want to attach this transistor to a path in a way that the transistor induces a step wise change of all the nodes in the path. Note that $T(k)$ contains only $\mathcal{O}(k)$ many vertices, yet its emitter vertices can potentially be connected to $\Omega(k^2)$ other vertices. Given a path graph of length $\mathcal{O}(k^2)$ and a transistor $T(k)$, the emitter vertices of the transistor are connected to the path in the following way: The first vertex in the path is connected to exactly two emitter vertices, the last is connected to none and each of the remaining nodes of the path is connected to exactly one emitter vertex. Furthermore, the collector edges of transistors of opinion X are always connected to constant reservoirs of opinion $-X$. Such a reservoir can be implemented as a clique. An illustration of this graph with $k = 3$ is given in Figure 3.4. Without loss of generality, we set the




Figure 3.4: Path with 4 vertices connected to one transistor $T(3)$.




Figure 3.5: Path with 4 vertices connected to 3 transistors $T(3)$. Note that transistors at bottom of figures are always upside down.




Figure 3.6: Two copies of Figure 3.5 with inverse opinions




Figure 3.7: In this graph, every time the path has run through completely the next transistor will flip, causing the path to run again.




Figure 3.8: Final graph in which the paths run 3 times.

initial state of the nodes of the path to B , and that of the transistor to R . As long as the transistor remains red, the path will turn red one vertex at a time. As soon as the transistor flips its opinion to blue (as a result of both base edges having advertised blue) the path will turn blue again, one vertex at a time. To force the path to change k times, k transistors are needed. Each of these transistors (note that we make use of red as well as blue transistors) is connected with the path in the same way as the first transistor. The resulting graph is given in Figure 3.5. A series of k switches of the complete path can now be provoked by switching transistors of alternating opinions in turns. For the example depicted in the figures, the switching order of the transistors is given by their respective indices.

Now, a way is needed to flip the next transistor every time the last vertex of the path has changed its opinion. Assume the last vertex has changed to red. It is necessary to flip a red transistor to blue in order to change the path to blue; however, the path changing to red can only cause a blue transistor to turn red. To this end, the graph is extended by a copy of itself with all opinions inverted. The resulting graph is given in Figure 3.6. As in every round each vertex in the copy has the opposite opinion of the original vertex, the copy of the last vertex in the path enables us to flip a red transistor to blue as desired. The edges necessary to achieve this (highlighted in green in Figure 3.7) connect the end of a path to \mathcal{B}^1 of each transistor in the other half of the graph. To ensure that the transistors flip in the required order, additional edges (highlighted in magenta in Figure 3.7) are introduced, connecting an emitter node of each transistor T_i^X to the node \mathcal{B}^1 of transistor T_{i+1}^X .

The green edges cause an unwanted influence on the last vertex of the paths. This influence can be negated by introducing additional edges (highlighted in cyan in Figure 3.8). These edges connect the last vertex of each path with an emitter vertex of each transistor not yet connected to that vertex.

The resulting graph contains $\mathcal{O}(k^2)$ vertices, yet has a convergence time of $\Omega(k^3)$. In terms of the number of vertices n , the convergence time is $n^{3/2}$. The detailed proof in Section 3.4.3 shows that this technique can be applied to run the entire graph repeatedly, just as the graph in this section runs two paths repeatedly, which leads to a convergence time

p	n	c
1	10	1
2	12	2
3	96	22
10	494	310
20	1614	3331
30	2010	5701
100	5518	45985

Figure 3.9: Overview of the exact convergence time c for some graphs from our graph class, constructed with n nodes and path length p .

of $\Omega(n^{7/4})$. In this new graph, the transistors change back and fourth repeatedly, always changing to the opinion advertised over the collector edges, similar to real transistors. When applied recursively $\log \log n$ times, an asymptotic convergence time of $\Omega(n^2/(\log \log n)^2)$ is reached. As the full proof is long and involved, we complement the formal proof with simulating these recursively constructed networks for path lengths of up to 100. Table 3.9 and Figure 3.10 show the outcomes of these simulations.

3.4.1 Transistor

In this section, we formally define how a transistor must be connected to the rest of the graph and how it behaves in that case. The symbol T is used to denote a particular graph as well as instances of that graph, which are induced subgraphs; the symbols $\mathcal{C}(T), \mathcal{E}(T), \mathcal{B}^x(T)$ are used to denote the respective vertex sets of an instance T . In order to talk about how a transistor should fit in a network, we introduce the outside influence function, which specifies how an induced subgraph is influenced by the rest of the graph.

Definition 3.5. For any vertex $v \in V^H$, where $H = (V^H, E^H)$ is an induced subgraph of $G = (V, E)$, we define an outside influence at time t exerted on said vertex as: $I_t^H(v) = \mu_0(v) \cdot \sum_{u \in \{u' | \{v, u'\} \in E \setminus E^H\}} \mu_t(u)$.

As all the nodes inside the same transistor start with the same opinion, the influence onto each other will be towards staying with this opinion.




Figure 3.10: Here one can see the simulation results compared to a quadratic curve. The point clusters arise when for several consecutive path lengths no new transistor is created. Small jumps in the number of vertices indicate that a new transistor was added; big jumps indicate that a new layer of transistors was added.

The only reason nodes change their opinion is, if the influence from outside the transistor makes them to. So if $I_t^H(v)$ is positive, node v will be influenced by the rest of the graph to stick with its initial opinion or change back to it; if $I_t^H(v)$ is negative, v will be influenced towards changing away from its initial opinion. The upper right indices are sometimes left out if they are clear from the context. We are now able to formally specify an outside influence range in which the transistor will operate correctly.

Definition 3.6. An instance $T = (V^T, E^T)$ of a transistor $T(k)$ with $k \geq 1$ and with outside influence $I_t^T(\cdot)$ is *correctly accessed* with initial opinion X if and only if the following conditions hold.

- (i) $\mu_0(v) = X$ for all $v \in V^T$
- (ii) $I_t(v) = -k$ for all t and all $v \in \mathcal{C}(T)$
- (iii) $|I_t(v)| \leq k - 1$ for all t and all $v \in \mathcal{E}(T)$
- (iv) $I_t(\mathcal{B}^1(T)) \in \{-3, -1, 1\}$ for all t
- (v) $I_t(\mathcal{B}^2(T)) = -1$ for all t
- (vi) $I_t(\mathcal{B}^3(T)) = -(k + 1)$ for all t

Note that if all collector edges in Figure 3.3 advertise R , the conditions (ii) through (vi) are fulfilled independent of what other edges connecting the transistor to the rest of the graph advertise. Because of (ii) there is a strong outside influence on the collector vertices to change their opinion. The very first round t where $I_t(\mathcal{B}^1(T)) = -3$, (this means both base edges in Figure 3.3 advertise R), will force all vertices of this transistor to flip their opinion to $(-X)$ eventually. We call this event t flip time $t^f(T)$ of T .

Lemma 3.7. *If an instance $T = (V^T, E^T)$ of a transistor $T(k)$ is correctly accessed with initial opinion X , then the following statements hold.*

- (i) All vertices $v \in V^T$ have opinion $\mu_t(v) = X$ for all $t \leq t^f(T)$
- (ii) All vertices $v \in \mathcal{E}(T)$ have opinion $\mu_t(v) = X$ for all $t^f(T) < t \leq t^f(T) + 3$

(iii) All vertices $v \in V^T$ have opinion $\mu_t(v) = (-X)$ for all $t \geq t^f(T) + 4$

Proof. First note that the vertices in $\mathcal{E}(T)$ can not change their opinion until at least one in $\mathcal{C}(T)$ has done so since they each have an outside influence of at most $(k - 1)$ times $(-X)$ after (iii) of definition 3.6 and an inside influence of k times X from the k vertices in $\mathcal{C}(T)$. Similarly, the vertices in $\mathcal{C}(T)$ can not change their opinion until $\mathcal{B}^3(T)$ and $\mathcal{B}^2(T)$ have done so and they in turn have to wait for $\mathcal{B}^1(T)$ to change. Finally, $\mathcal{B}^1(T)$ will only change after the flip time $t^f(T)$. This induces (i) and (ii). To see that (iii) is true note that after the flip event the sets $(\{\mathcal{B}^1(T)\}, \{\mathcal{B}^2(T), \mathcal{B}^3(T)\}, \mathcal{C}(T), \mathcal{E}(T))$ will indeed all change their opinion in the given order. Also note that even if the outside influence of $\mathcal{B}^1(T)$ goes back to some number $x > -3$, the process started by the flip event can not be stopped or reversed. \square

3.4.2 Counters

The final graph for our intended lower bound is a recursively defined counter. In this section, we present the base case in form of a simple 2-Path graph as well as the recursive case. In the latter, a counter is combined with a number of transistors to form a bigger counter as suggested in the proof outline.

A counter $K = (H = (V, E), \mathcal{I}(\cdot), \mathcal{R}^R, \mathcal{R}^B, \mathcal{S}(\cdot))$ consists of a graph $H = (V, E)$, a function $\mathcal{I} : V \rightarrow \mathbb{Z}$, specifies the valid range of outside influence, two special interest vertices $\mathcal{R}^R, \mathcal{R}^B \in V$, which indicate when the graph has finished running, and an initial configuration $\mathcal{S}(\cdot)$.

We will postpone the definition of the axioms a counter must satisfy, and first describe how a counter is properly connected and accessed since the behavior of a counter only needs to be defined if it is connected and accessed correctly.

Definition 3.8. A counter $K = (H = (V^H, E^H), \mathcal{I}(\cdot), \mathcal{R}^R, \mathcal{R}^B, \mathcal{S}(\cdot))$ is *correctly accessed* and *correctly initialized* from t_1 to t_2 if and only if the following condition holds.

- (i) For all v in V^H the initial state of v is set by $\mu_{t_1}(v) = \mathcal{S}(v)$.

- (ii) For all X and all $t_1 \leq t \leq t_2$ the outside influence is given by $I_t^H(v) = \mathcal{I}(v)$

We say a counter K is *reversely initialized* if definition (i) is changed to $\mu_{t_1}(v) = -\mathcal{S}(v)$, and the counter is considered *reversely accessed* if (ii) is changed to $I_t^H(v) = -\mathcal{I}(v)$. We sometimes add the keyword *virtually* to indicate some deviations from the definition which do not result in an altered behavior of H .

Definition 3.9. A tuple $K = (H = (V^H, E^H), \mathcal{I}(\cdot), \mathcal{R}^R, \mathcal{R}^B, \mathcal{S}(\cdot))$ is a proper counter with convergence time c and supply edge number e if and only if a correct access and a correct initialization from t_1 to t_2 imply that the following statements hold.

- (i) $e = \sum_{\{v \in V^K \mid \mathcal{I}(v) \cdot \mathcal{S}(v) > 0\}} \mathcal{I}(v) \cdot \mathcal{S}(v) = \sum_{\{v \in V^K \mid \mathcal{I}(v) \cdot \mathcal{S}(v) < 0\}} -\mathcal{I}(v) \cdot \mathcal{S}(v)$
- (ii) $\lceil \sqrt{e+2} \rceil + 1 \geq \max_{v \in V^K} |\mathcal{I}(v)|$
- (iii) The vertices \mathcal{R}^X are of opinion X for all $t \leq \min\{t_1 + c - 1, t_2 + 1\}$
- (iv) For all vertices $v \in V^H$ and all $t_1 + c \leq t \leq t_2 + 1$ the following is true $\mu_t(v) = -\mu_{t_1}(v)$.

The edge supply number corresponds to the number of black edges in Figure 3.4 and (i) satisfied if it is the same for blue and red. We will need condition (ii) to make sure that we do not produce a multigraph, (iii) indicates that \mathcal{R}^X only change their opinion after l rounds when the graph has finished running, and (iv) makes it possible to run the counter again with reverse opinions. Note that if a counter is correctly initialized and reversely accessed, or if it is reversely initialized and correctly accessed from t_1 to t_2 , $\mu_t(v)$ will be time-constant from t_1 to t_2 because of (iv).

2-Path graph The 2-Path graph counter will be the base case of our final, recursively defined graph. It consists just of two simple paths one of which is initialized with R and the other with B . If accessed correctly, the following will happen on both paths simultaneously. All vertices will change their opinion in the order in which they occur on their respective paths.

Definition 3.10. In the following we define a 2-Path graph of length l as $P_2(l) = (H = (V^H, E^H), \mathcal{I}(\cdot), \mathcal{R}^R, \mathcal{R}^B, \mathcal{S}(\cdot))$, where

$$\begin{aligned} V^H &= \{v_i^R \mid 0 \leq i < l\} \cup \{v_i^B \mid 0 \leq i < l\} \\ E^H &= \{\{v_i^R, v_{i+1}^R\} \mid 0 \leq i < l-1\} \cup \{\{v_i^B, v_{i+1}^B\} \mid 0 \leq i < l-1\} \\ \mathcal{I}(v) &= \begin{cases} -2 & \text{if } v = v_0^X \\ -1 & \text{if } v \in \{v_i^X \mid 1 \leq i < l-1\} \\ 0 & \text{otherwise} \end{cases} \\ \mathcal{R}^X &= v_{l-1}^X \\ \mathcal{S}(v_i^X) &= X \end{aligned}$$

Lemma 3.11. A 2-Path graph $P_2(l)$ is a valid counter with $n = 2l$ vertices, convergence time $c = l$ and supply edge number $e = l$.

Proof. We have to prove the conditions in definition 3.9 under the assumption that the conditions in the definition 3.8 hold.

The conditions (i) and (ii) from Definition 3.9 are trivially true. For (iii) note that from condition (ii) of definition 3.8 and the definition of $\mathcal{I}(\cdot)$, it follows that the vertices in $\{v_i^X \mid 1 \leq i < l-1\}$ all have one more outside neighbor of the opinion $-X$ than of the opinion X , the vertices v_0^X has two more outside neighbors of the opinion $(-X)$ than of the opinion X , and v_{l-1}^X has the same number of outside neighbors of the opinion $(-X)$ as of the opinion X . This means that the vertices v_0^X will change its opinion to $(-X)$ in the first round, and cause v_1^X to turn in the second and so forth. In other words, $\mu_t(v_i^X)$ will be X for all $t \leq i$ and $(-X)$ ever after. Therefore, $\mu_t(\mathcal{R}^X) = \mu_t(v_{l-1}^X)$ is equal to X for all $t \leq l-1$, and condition (iii) of definition 3.9 is satisfied.

Also note that at time l all vertices have changed their opinion and are not going to reverse back to their original opinion therefore satisfying condition (iv). \square

Repeater A repeater is a function that takes a counter and uses transistors to repeatedly run that counter in order to produce a counter with much higher convergence time at the expense of an only slightly increased

number of vertices. However in addition to what was suggested earlier in this section, we need two new vertices $\mathcal{R}^R, \mathcal{R}^B$ to indicate when the new graph has reached a stable state as displayed in Figure 3.11.

Definition 3.12. A repeater is a function R which when it is given a number j and a counter $\tilde{K} = (\tilde{H} = (\tilde{V}^H, \tilde{E}^H), \tilde{\mathcal{I}}(\cdot), \tilde{\mathcal{R}}^R, \tilde{\mathcal{R}}^B, \tilde{\mathcal{S}}(\cdot))$ with convergence time \tilde{c} , \tilde{n} vertices, and supply edge number \tilde{e} , produces a tuple $R(\tilde{K}, j) = (H = (V^H, E^H), \mathcal{I}(\cdot), \mathcal{R}^R, \mathcal{R}^B, \mathcal{S}(\cdot))$ where

$$T_i^X = (V_i^X, E_i^X) \text{ for } 0 \leq i \leq 2j, X \in \{R, B\},$$

are instances of transistors $T(s)$ and

$$\begin{aligned} V^H &= \left(\bigcup_{X \in \{R, B\}, i=0}^{j-1} V_i^X \right) \cup \tilde{V}^H \cup v^R \cup v^B \\ E^H &= \left(\bigcup_{X \in \{R, B\}, i=0}^{j-1} E_i^X \right) \cup \tilde{E}^H \cup E^1 \cup E^2 \cup E^3 \cup E^4 \cup E^5 \\ E^1 &= \{ \{ \tilde{\mathcal{R}}^X, \mathcal{B}^1(T_{2i}^{-X}) \} \mid 0 \leq i \leq 2j \} \\ &\quad \cup \{ \{ \tilde{\mathcal{R}}^X, \mathcal{B}^1(T_{2i+1}^X) \} \mid 0 \leq i \leq 2j - 1 \} \text{ (green in Figure 3.11)} \\ E^2 &= \{ \{ \tilde{\mathcal{R}}^X, \mathcal{R}^X \} \mid X \in \{R, B\} \} \text{ (orange in Figure 3.11)} \end{aligned}$$

$$\mathcal{I}(v) = \begin{cases} -2 & \text{if } v = \mathcal{B}^1(T_0^X) \\ -1 & \text{if } v = \mathcal{B}^1(T_i^X) \text{ for some } 1 \leq i \leq 2j \\ -1 & \text{if } v = \mathcal{B}^2(T_i^X) \text{ for some } 0 \leq i \leq 2j \\ -(s+1) & \text{if } v = \mathcal{B}^3(T_i^X) \text{ for some } 0 \leq i \leq 2j \\ -s & \text{if } v \in \mathcal{E}(T_i^X) \\ 1 & \text{if } v = \mathcal{R}^X \\ -1 & \text{if } v = \tilde{\mathcal{R}}^X \\ 0 & \text{otherwise} \end{cases}$$

$$\mathcal{R}^R = v^R$$

$$\mathcal{R}^B = v^B$$

$$\mathcal{S}(v) = \begin{cases} \tilde{\mathcal{S}}(v) & \text{if } v \in \tilde{V}^T \\ X & \text{if } v \in V_i^X \\ X & \text{if } v = \mathcal{R}^X \end{cases}$$

$$s = \lceil \sqrt{\tilde{e} + 2} \rceil + 1$$

The edge set E^3 (black in Figure 3.11) consists of the following edges. For every vertex $v \in \tilde{V}^H$ with $k = |\tilde{\mathcal{I}}(v)|$ and with $X = \text{sign}(\tilde{\mathcal{I}}(v)\mathcal{S}(v))$, there are edges from v to k different emitter vertices of every T_{2i}^{-X} with $0 \leq i \leq j$ and of every T_{2i+1}^X with $0 \leq i \leq j-1$.

The edge set E^4 (magenta in Figure 3.11) consists of the following edges. For every $0 \geq i \geq 2j-1$ and X , there is an edge from $\mathcal{B}^1(T_{i+1}^X)$ to an emitter vertex of T_i^X and there is one edge from an emitter vertex of T_{2j}^X to \mathcal{R}^X .

The edge set E^5 (cyan in Figure 3.11) consists of the following edges. For all X , there is an edge from $\tilde{\mathcal{R}}^X$ to an emitter vertex of every T_{2i}^X with $0 \leq i \leq j$ and of every T_{2i+1}^{-X} with $0 \leq i \leq j-1$.

Lemma 3.13. *The repeater $R(\tilde{K}, j)$ does exist in a way that the emitter vertices of T_i are connected to no more than $(s-1)$ vertices outside T_i .*

Proof. We have to show that the edge sets E^3 , E^4 and E^5 exist, such that the emitter vertices of the transistors do not have too many outside edges. Every transistor's emitter vertices have \tilde{e} edges leaving the transistor in




Figure 3.11: Full repeater graph $R(1, P_2(4))$ with a $P_2(4)$ as graph which is repeatedly run and six transistors to run $P_2(4)$ three times.

E^3 because of condition (i) of definition 3.8, one edge in E^4 and one edge in E^5 . This is a total of $\tilde{e} + 2$ necessary edges. A transistor $T(s)$ has s collector vertices each of which should have no more than $(s - 1)$ outside edges where $s = \lceil \sqrt{\tilde{e} + 2} \rceil + 1$. This makes a total of more than possible $\tilde{e} + 2$ edges.

$$\begin{aligned} s(s - 1) &= (\lceil \sqrt{\tilde{e} + 2} \rceil + 1)(\lceil \sqrt{\tilde{e} + 2} \rceil + 1 - 1) \\ &> \lceil \sqrt{\tilde{e} + 2} \rceil^2 \\ &\geq \tilde{e} + 2 \end{aligned}$$

Additionally condition (ii) of definition 3.8 shows that $|\tilde{\mathcal{I}}(v)| \leq s$. This is necessary because we could otherwise only realize the graph if we were allowed multigraphs. \square

Lemma 3.14. *The repeater $R(\tilde{K}, j)$ has $n = 2(2j + 1)(2s + 1) + \tilde{n} + 2$ vertices and supply edge number $e = (2j + 1)(s^2 + s + 3) + 3$ where $s = \lceil \sqrt{\tilde{e} + 2} \rceil + 1$.*

Proof. To obtain the number of vertices we add the number of vertices in a transistor times the number of transistors to the number in the counter K and two vertices \mathcal{R}^X .

$$\begin{aligned} |V^H| &= 2(2j + 1)|T(s)| + |\tilde{V}^H| + 2 \\ &= 2(2j + 1)(2s + 1) + n + 2 \end{aligned}$$

To obtain the supply edge number we just go through the cases in the definition of $\mathcal{S}(\cdot)$ and add them up.

$$\begin{aligned} e &= 1 \cdot 2 + 2j + (2j + 1) + (2j + 1) \cdot (s + 1) + (2j + 1)s^2 + 1 \cdot 1 + 1 \cdot 1 \\ &= 2(2j + 1) + (2j + 1)(s + 1) + (2j + 1)s^2 + 3 \\ &= (2j + 1)(s^2 + s + 3) + 3 \end{aligned}$$

\square

Lemma 3.15. *For all $0 \leq i \leq 2j$ and all X , T_i^X is accessed correctly.*

Proof. Condition (i) of Definition 3.6 is trivially true. The collector vertices as well as \mathcal{B}^2 and \mathcal{B}^3 of T_i have no edges to other parts of H so it holds that $I_t^{T_i}(v) = I_t^H(v) = \mathcal{I}(v)$ for all $v \in \mathcal{C}(T_i) \cup \{\mathcal{B}^2(T_i)\} \cup \{\mathcal{B}^3(T_i)\}$. Therefore, (ii),(v) and (vi) are fulfilled.

Condition (iii) is fulfilled because of Lemma 3.13.

For condition (iv) we distinguish between $i = 0$ and $i \neq 0$. If i is equal to 0, then $\mathcal{B}^1(T_i^X)$ has 1 edge coming from outside V_i^X but still inside H and an outside influence of $I^H(\mathcal{B}^1(T_i^X)) = -2$. If i is not equal to 0, then $\mathcal{B}^1(T_i^X)$ has 2 edges coming from outside V_i^X but still inside H and an outside influence of $I^H(\mathcal{B}^1(T_i^X)) = -1$. In both cases, the resulting $I^{T_i^X}(\mathcal{B}^1(T_i^X))$ will satisfy (iv). \square

Lemma 3.16. *The transistors flip in order. That is, the following statements must hold.*

(i) *If $t^f(T_i^X) = t$ there must be a $t' \leq t - 4$ such that $t^f(T_{i-1}^X) = t'$ for all $1 \leq i \leq 2j$.*

(ii) *If $O_t^{\mathcal{R}^X}(\mathcal{R}^X) = -X$ there must be a $t' \leq t - 4$ such that $t^f(T_{2j}^X) = t'$*

Proof. The vertex \mathcal{R}^X has total influence $I_t^v(\mathcal{R}^X) = I_t^H(\mathcal{R}^X) + \mu_t(\tilde{\mathcal{R}}^X) + \mu_t(\mathcal{E}(T_{2j}^X))$ which is equal to $1 + X \cdot (\mu_t(\tilde{\mathcal{R}}^X) + \mu_t(\mathcal{E}(T_{2j}^X)))$. So \mathcal{R}^X can only change its opinion to $(-X)$ if both $\mu_t(\tilde{\mathcal{R}}^X)$ and $\mu_t(\mathcal{E}(T_{2j}^X))$ are $(-X)$. Similarly $I_t^{T_i^X}(\mathcal{B}(T_i^X))$ which can be written as $I_t^H(\mathcal{B}(T_i^X)) + X(\mu_t(\mathcal{E}(T_{i-1}^X)) + \mu_t(\tilde{\mathcal{R}})) = -1 + X(\mu_t(\mathcal{E}(T_{i-1}^X)) + \mu_t(\tilde{\mathcal{R}}))$ can only be -3 if $\mathcal{E}(T_{i-1}^X)$ is $-X$ for all $1 \leq i \leq 2j$. This together with statements (i) and (ii) of Lemma 3.7 we get the required statements. \square

Definition 3.17. We define ${}^1I^{\tilde{H}}(\cdot)$, ${}^2I^{\tilde{H}}(\cdot)$, ${}^3I^{\tilde{H}}(\cdot)$ and ${}^5I^{\tilde{H}}(\cdot)$ to be the outside influence exerted on vertices in \tilde{H} by E^1 , E^2 , E^3 and E^5 respectively.

Note that $I_t^{\tilde{H}}(v) = {}^1I_t^{\tilde{H}}(v) + {}^2I_t^{\tilde{H}}(v) + {}^3I_t^{\tilde{H}}(v) + {}^5I_t^{\tilde{H}}(v) + I_t^H(v)$. The edges E^4 are intentionally left out since they do not have endpoints in \tilde{H} .

Lemma 3.18. *$R(j, \tilde{K})$ is indeed a counter of convergence time $c = (2j + 1)(\tilde{c} + 4) + 1$.*

Proof. So assume that the constant outside influence of all vertices is given by $I^H(\cdot) = \mathcal{I}^H(\cdot)$. We define c to be the smallest t with $\mu_t(\mathcal{R}^R) = B$ or $\mu_t(\mathcal{R}^B) = R$. By Lemma 3.16 all transistors have flipped at time c . For all $t < c$ and all vertices $v \in \tilde{H}$, the outside influences of ${}^2I_t^{\tilde{H}}(v)$ and of $I_t^H(v)$ cancel each other out so we only need to care about ${}^1I^{\tilde{H}}()$, ${}^3I^{\tilde{H}}()$ and ${}^5I^{\tilde{H}}()$ until we reach time c .

To get a induction hypothesis we prove the following stronger statement: $t^f(T^{2l}) = 2l(\tilde{c} + 4)$ and $\mu_{2l(\tilde{c}+4)}(v) = \mathcal{S}(v)$ for all $0 \leq l \leq j$ and for all vertices v in \tilde{H} .

The base case for $l = 0$ is trivial. For the induction step $l \rightarrow l + 1$ assume $t^f(T^{2l}) = 2l(\tilde{c} + 4)$ and $\mu_{2l}(v) = \mathcal{S}(v)$ is true for all vertices v in \tilde{H} . Now we investigate how the graph evolves in the following intervals $[2l(\tilde{c} + 4), 2l(\tilde{c} + 4) + 3]$, $[2l(\tilde{c} + 4) + 4, (2l + 1)(\tilde{c} + 4) - 1]$, $[(2l + 1)(\tilde{c} + 4), (2l + 1)(\tilde{c} + 4) + 3]$ and $[(2l + 1)(\tilde{c} + 4) + 4, (2l + 2)(\tilde{c} + 4) - 1]$. I.e., we take the behavior during the interval $[a, b]$ to investigate the influences exerted in $[a, b]$ and their outcomes in $[a + 1, b + 1]$.

Interval $[2l(\tilde{c} + 4), 2l(\tilde{c} + 4) + 3]$. We show that \tilde{K} is correctly initialized and virtually reversely accessed from $2l(\tilde{c} + 4)$ to $2l(\tilde{c} + 4) + 3$. Therefore and because of (iii) in definition 3.9, $\mu_{2l(\tilde{c}+4)+4}(\cdot)$ will still be equal to $\mathcal{S}(\cdot)$.

Let us start by considering ${}^1I^{\tilde{H}}(\cdot) + {}^5I^{\tilde{H}}(\cdot)$. Using Lemma 3.16 we can deduce the following statement. All transistors T_i^X with $i < 2l$ have already switched completely by $2l(\tilde{c} + 4)$ (they are in the region specified by (iii) of Lemma 3.7, similarly all such transistors with $i > 2l$ can only start switching after $i(\tilde{c} + 4) + 3$ (they are in the region specified by (i) of Lemma 3.7. So the contribution of T_i^Y to ${}^1I^{\tilde{H}}()$ is canceled out by the contribution of T_i^{-Y} to ${}^5I^{\tilde{H}}()$ for all $i \neq 2l$ and vice versa. And for T_{2l}^X we know that $\mu_i(\mathcal{E}(T_{2j}^X)) = X$ for $i \leq 2l(\tilde{c}_4) + 3$ from (i), (ii) of Lemma 3.7. Therefore, ${}^1I_i^{\tilde{H}}(v) + {}^5I_i^{\tilde{H}}(v)$ must be non-negative for $\tilde{\mathcal{R}}^X$ and zero for all other vertices.

${}^3I_i^{\tilde{H}}(\cdot)$ will be exactly $-\mathcal{I}(\cdot)$ for $2l(\tilde{c} + 4) \leq i \leq 2l(\tilde{c} + 4) + 3$ because the transistors T_i^Y with an even $i \neq 2l$ cancel out each other and those with $i = 2l$ will have $\mu_t(\mathcal{E}(T_i^X)) = X$ for $t \leq 2l(\tilde{c} + 4) + 3$ so T_i^X will exactly contribute the required outside influence of $-\mathcal{I}(\cdot)$. So $I_t^T(v) =$

${}^1I_t^{\tilde{H}}(v) + {}^3I_t^{\tilde{H}}(v) + {}^5I_t^{\tilde{H}}(v)$ will be $-\mathcal{I}(v)$ for all v except for \mathcal{R}^R and \mathcal{R}^B . However \mathcal{R}^X is supposed to stick with $o(\mathcal{R}^X) = X$, in case of correct initialization and reverse access and the deviation introduced by ${}^1I_t^{\tilde{H}}(\cdot) + {}^5I_t^{\tilde{H}}(\cdot)$ further encourages them to do so. Therefore \tilde{K} is a correctly initialized and virtually reversely accessed counter.

Interval $[2l(\tilde{c}+4)+4, (2l+1)(\tilde{c}+4)-1]$. We will use a proof of induction over k to show that \tilde{K} is accessed and initialized correctly from $2l(\tilde{c}+4)+4$ to $2l(\tilde{c}+4)+4+k$ and that $t^f(T_{2l+1}^X) \geq 2l(\tilde{c}+4)+k$ for all $k < \tilde{c}$. Using (iv) of definition 3.9 the induction statement for k set to $\tilde{c}-1$ will imply that $\mu_{(2l+1)(\tilde{c}+4)}(\cdot) = -\tilde{S}(\cdot)$. This in turn will imply that $t^f(T_{2l+1}^X) = (2l+1)(\tilde{c}+4)$.

Because of Lemma 3.16 and because $\mu_{2l(\tilde{c}+4)}(\tilde{\mathcal{R}}^X) = X$, $t^f(T_{2l+1}^e)$ must be bigger than $2l(\tilde{c}+4)+4$. Therefore, each transistor has a uniform opinion at $2l(\tilde{c}+4)$, and therefore ${}^1I_{2l+4}^{\tilde{H}}(\cdot) + {}^5I_{2l+4}^{\tilde{H}}(\cdot)$ will be 0. ${}^3I_{2l+4}^{\tilde{H}}(v)$ will be exactly $\mathcal{I}(v)$. This is because of the same reason as in the previous interval. All transistors T_i^X with $i \neq 2l$ cancel out the transistors $i = 2l$ that have already flipped completely. Hence, \tilde{K} is accessed and initialized correctly from $2l(\tilde{c}+4)+4$ to $2l(\tilde{c}+4)+4$. This covers the base case of our induction.

Now assuming \tilde{K} is accessed and initialized correctly from $2l(\tilde{c}+4)+4$ to $2l(\tilde{c}+4)+4+k$ and $t^f(T_{2l+1}^X) \geq 2l(\tilde{c}+4)+k$ for some $k < \tilde{c}$. Because of the first assumption and because $k+1 < \tilde{c}$ we can apply (iii) of definition 3.9. Hence $\mu_{2l(\tilde{c}+4)+k}(\tilde{\mathcal{R}}^X) = X$ will still be true and as a direct consequence $t^f(T_{2l+1}^X) \geq 2l(\tilde{c}+4)+k+1$ and we have proven the first part of our induction statement. Now since $t^f(T_{2l+1}^X) \geq 2l(\tilde{c}+4)+k+1$ T_{2l+1}^X is still uniformly of opinion X at $2l(\tilde{c}+4)+k+1$, so $I_{2l(\tilde{c}+4)+k+1}^{\tilde{H}}(\cdot)$ is still equal to $\mathcal{I}(\cdot)$. This also means that \tilde{K} is still correctly accessed and therefore proves our induction statement.

Interval $[(2l+1)(\tilde{c}+4), (2l+1)(\tilde{c}+4)+3]$ The graph \tilde{K} is reversely initialized and virtually correctly accessed from $(2l+1)(\tilde{c}+4)$ to $(2l+1)(\tilde{c}+4)+3$. We know from the last interval that \tilde{K} is reversely initialized and that $T_{(2l+1)}^X = (2l+1)(\tilde{c}+4)$ The proof of correct access is completely analogous to the first interval. Therefore and because of (iii) in definition

3.9, $\mu_{2l(\tilde{c}+4)+4}(\cdot)$ will still be equal to $-\mathcal{S}(\cdot)$.

Interval $[(2l+1)(\tilde{c}+4)+4, (2l+2)(\tilde{c}+4)-1]$ The graph \tilde{K} is accessed and initialized reversely from $(2l+1)(\tilde{c}+4)+4$ to $(2l+2)(\tilde{c}+4)-1$ and $t^f(I_{2l+2}^X) = 2l(\tilde{c}+4)$. We know from the last interval that \tilde{K} is reversely initialized. The rest of the proof is completely analogous to the second interval. Therefore $\mu_{(2l+2)(\tilde{c}+4)}(\cdot) = \tilde{\mathcal{S}}(\cdot)$.

This concludes the proof of the induction step. Now we only need to show that we can indeed derive the lemma from this induction result. At time $2j(\tilde{c}+4)$ we basically have the same case as at the beginning of the first interval and using the same technique as in the first and second interval we can therefore deduce that the smallest time t with $I_t^{\tilde{H}}(\mathcal{R}^X) < 0$ is $(2j+1)(\tilde{c}+4)$ (this proves (iii) of definition 3.9). We can in the same manner prove that $\mu_{(2j+1)(\tilde{c}+4)}(v) = -\mathcal{S}(v)$ for all $v \in \tilde{V}^H$. And since all transistors have already flipped completely and because \tilde{K} is reversely initialized and correctly accessed from $(2j+1)(\tilde{c}+4)$ to infinity, (iv) is also fulfilled. The conditions (i) are true because of Lemma 3.14 and finally (ii) is true since

$$\begin{aligned} \lceil \sqrt{e+2} \rceil + 1 &= \lceil \sqrt{(2j+1)(s^2+s+3)+3+2} \rceil + 1 \\ &\geq \lceil \sqrt{s^2} \rceil + 1 \\ &= s + 1 \\ &= \max_{v \in V^H} \mathcal{I}(v). \end{aligned}$$

□

3.4.3 Putting it all together

Lemma 3.19. *For every $l \geq 1$ and $h \geq 0$, there exists a counter K with $n \leq (h+1)54 \left(1 + \frac{80}{\sqrt{l}}\right)^h l$ vertices, supply edge number $e \leq l^{2-\frac{1}{2^h}}$. $\left(1 + \frac{80}{\sqrt{l}}\right)^h$ and convergence time $c \geq l^{2-\frac{1}{2^h}}$.*

Proof. We prove this by induction over h . For $h=0$, K is trivially given by $P_2(l)$. Now given a counter \tilde{K} with $\tilde{n} \leq (h+1)54 \left(1 + \frac{80}{\sqrt{l}}\right)^h l$ vertices,

convergence time $\bar{c} \geq l^{2-\frac{1}{2^h}}$ and supply edge number

$$\bar{e} \leq l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h$$

we can construct

$$K = R(\tilde{K}, \frac{1}{2} \lfloor l^{\frac{1}{2^{h+1}}} \rfloor).$$

By Lemma 3.18 the convergence time c is at least $(2\frac{1}{2} \lfloor l^{\frac{1}{2^{h+1}}} \rfloor + 1)(l^{2-\frac{1}{2^h}} + 4)$ which is in turn at least $l^{2-\frac{1}{2^{h+1}}}$. To prove the bound on the supply edge number, we first show a bound for the transistor size s by using the definition of R . The required bound for the supply edge number of K can be deduced using Lemma 3.14.

$$s = \lceil \sqrt{\bar{e} + 2} \rceil + 1 \tag{3.1}$$

$$\leq \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} + 2 + 2 \tag{3.2}$$

$$\leq \sqrt{\left(1 + \frac{2}{l}\right) l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} + 2 \tag{3.3}$$

$$\leq \left(1 + \frac{2}{l}\right) \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} + 2 \tag{3.4}$$

$$\leq \left(1 + \frac{2}{l} + \frac{2}{\sqrt{l}}\right) \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} \tag{3.5}$$

$$\leq \left(1 + \frac{4}{\sqrt{l}}\right) \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} \tag{3.6}$$

$$\begin{aligned}
e &= (2j+1)(s^2+s+3)+3 && \text{where } j = \frac{1}{2} \lfloor l^{\frac{1}{2^{h+1}}} \rfloor \\
&\leq \left(1 + \frac{1}{l}\right) \cdot l^{\frac{1}{2^{h+1}}} (s^2+s+3)+3 \\
&\leq \left(1 + \frac{1}{l}\right) \cdot l^{\frac{1}{2^{h+1}}} \left(\left(s + \frac{1}{2}\right)^2 + \frac{11}{4} \right) + 3 \\
&\leq \left(1 + \frac{1}{l}\right) \cdot l^{\frac{1}{2^{h+1}}} \left(\left(\left(1 + \frac{4}{\sqrt{l}}\right) \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} + \frac{1}{2} \right)^2 + \frac{11}{4} \right) \\
&\quad + 3 \\
&\leq \left(1 + \frac{1}{l}\right) \cdot l^{\frac{1}{2^{h+1}}} \left(\left(\left(1 + \frac{5}{\sqrt{l}}\right) \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} + \frac{11}{4} \right)^2 + 3 \right) \\
&\leq \left(1 + \frac{1}{l}\right) \cdot l^{\frac{1}{2^{h+1}}} \left(\left(1 + \frac{10}{\sqrt{l}} + \frac{25}{l}\right) l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h + \frac{11}{4} \right) + 3 \\
&\leq \left(1 + \frac{1}{l}\right) \cdot l^{\frac{1}{2^{h+1}}} \left(1 + \frac{38}{\sqrt{l}}\right) l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h + 3 \\
&\leq \left(1 + \frac{77}{\sqrt{l}}\right) \cdot l^{2-\frac{1}{2^{h+1}}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h + 3 \\
&= \left(1 + \frac{80}{\sqrt{l}}\right)^{h+1} l^{2-\frac{1}{2^{h+1}}}
\end{aligned}$$

When contracting terms in (3.3) and (3.5), we utilize the fact that $l^{2-\frac{1}{2^h}} \geq l$ and $\left(1 + \frac{80}{\sqrt{l}}\right)^h \geq 1$. Using the same lemma, also gives the

needed bound for the number of vertices n .

$$\begin{aligned}
n &= 2(2j+1)(2s+3) + \tilde{n} + 2 && \text{where } \tilde{n} = (h+1)54 \left(1 + \frac{80}{\sqrt{l}}\right)^h l \\
&\leq 2 \left(1 + \frac{1}{l}\right) l^{\frac{1}{2^{h+1}}} (2s+3) + \tilde{n} + 2 \\
&\leq \left(2 + \frac{2}{l}\right) l^{\frac{1}{2^{h+1}}} \left(2 \left(1 + \frac{4}{\sqrt{l}}\right) \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} + 3\right) + \tilde{n} + 2 \\
&\leq \left(2 + \frac{2}{l}\right) l^{\frac{1}{2^{h+1}}} \left(2 + \frac{11}{\sqrt{l}}\right) \sqrt{l^{2-\frac{1}{2^h}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^h} + \tilde{n} + 2 \\
&\leq \left(4 + \frac{48}{\sqrt{l}}\right) l^{\frac{1}{2^{h+1}}} \cdot l^{1-\frac{1}{2^{h+1}}} \left(1 + \frac{80}{\sqrt{l}}\right)^{\frac{h}{2}} + \tilde{n} + 2 \\
&\leq \left(4 + \frac{50}{\sqrt{l}}\right) \left(1 + \frac{80}{\sqrt{l}}\right)^{\frac{h}{2}} l + \tilde{n} \\
&\leq 54 \left(1 + \frac{80}{\sqrt{l}}\right)^{\frac{h}{2}} l + (h+1)54 \left(1 + \frac{80}{\sqrt{l}}\right)^h l \\
&\leq (h+2)54 \left(1 + \frac{80}{\sqrt{l}}\right)^h l \\
&\leq (h+2)54 \left(1 + \frac{80}{\sqrt{l}}\right)^{h+1} l
\end{aligned}$$

□

We need one more additional tool from mathematics to prove our final Theorem 3.6.

Lemma 3.20. *For every $x > 0$ and all $n > 0$ the following inequality holds $\left(1 + \frac{x}{n}\right)^n \leq e^x$*

Proof. It is well known that the sequence $s_n(x) = \left(1 + \frac{x}{n}\right)^n$ converges to e^x as n goes to infinity. So we need only prove that $s_n(x)$ is non-decreasing in n . We achieve this by showing that all coefficients in the power series $s_n(x) = \sum_{k=0}^{\infty} c_n^k x^k = \sum_{k=0}^{\infty} \binom{n}{k} \left(\frac{x}{n}\right)^k$ are non decreasing.

$$\frac{c_{n+1}^k}{c_n^k} = \frac{\binom{n+1}{k} \frac{1}{(n+1)^k}}{\binom{n}{k} \frac{1}{n^k}} \quad (3.7)$$

$$= \frac{\frac{(n+1)!}{k!(n+1-k)!} \frac{1}{(n+1)^k}}{\frac{n!}{k!(n-k)!} \frac{1}{n^k}} \quad (3.8)$$

$$= \frac{\frac{1}{n+1-k} \frac{1}{(n+1)^{k-1}}}{\frac{1}{n^k}} \quad (3.9)$$

$$= \frac{n^k}{(n+1)^{k-1}(n-(k-1))} \quad (3.10)$$

$$\geq \frac{n^k}{n^k} \quad (3.11)$$

$$= 1 \quad (3.12)$$

In (3.10), the arithmetic mean of the factors in the denominator (numerator respectively) is n . Since the geometric mean of positive numbers can never be bigger than the arithmetic mean, the geometric mean of the factors of the denominator has to be $\leq n$ and the product has to be $\leq n^k$. \square

Now we combine Lemma 3.19 and Lemma 3.20 to prove Theorem 3.6.

Proof. If we select h in Lemma 3.19 to be $\lceil \log \log l \rceil$ we can get a counter

with the following dimensions for every l .

$$n \leq (\lfloor \log \log l \rfloor + 1) 54 \left(1 + \frac{80}{\sqrt{l}}\right)^{\lfloor \log \log l \rfloor} l \quad (3.13)$$

$$\leq 108 \log \log l \left(1 + \frac{80}{\lfloor \sqrt{l} \rfloor}\right)^{\lfloor \log \log l \rfloor} l \quad (3.14)$$

$$\leq 108 \log \log l \left(1 + \frac{80}{\lfloor \sqrt{l} \rfloor}\right)^{\lfloor \sqrt{l} \rfloor} l \quad (3.15)$$

$$\leq 108 \cdot e^{80} \log \log l \cdot l \quad (3.16)$$

$$e \leq l^2 \cdot l^{-\frac{1}{2\lfloor \log \log l \rfloor}} \cdot \left(1 + \frac{80}{\sqrt{l}}\right)^{\lfloor \log \log l \rfloor} \quad (3.17)$$

$$\leq l^2 \cdot l^{-\frac{1}{\log l}} \cdot e^{80} \quad (3.18)$$

$$= \frac{1}{2} l^2 \cdot e^{80} \quad (3.19)$$

$$c \geq l^2 \cdot l^{-\frac{1}{2\lfloor \log \log l \rfloor}} \quad (3.20)$$

$$\geq l^2 \cdot l^{-\frac{1}{2(\log \log l) - 1}} \quad (3.21)$$

$$= l^2 \cdot l^{-\frac{2}{\log l}} \quad (3.22)$$

$$= l^2 \cdot \left(l^{-\frac{1}{\log l}}\right)^2 \quad (3.23)$$

$$= \left(\frac{1}{2}\right)^2 l^2 \quad (3.24)$$

$$= \frac{1}{4} l^2 \quad (3.25)$$

(3.16) and (3.18) are true because of Lemma 3.20, and (3.19) and (3.24) are true because it holds that $\log\left(l^{-\frac{1}{\log l}}\right) = -\frac{1}{\log l} \log l = -1 = \log \frac{1}{2}$. We can also run this counter by creating a red and a blue clique of size $\lceil \sqrt{e} \rceil + 1$ and then connecting the vertices in the counter to vertices in the cliques according to $\mathcal{I}(\cdot)$. Since $\sqrt{e} = \mathcal{O}(n)$ this increases the number of vertices only by a constant fraction. This concludes that our

final network has $n = \mathcal{O}(l \cdot \log \log l)$ vertices and a convergence time of $\Omega(l^2) = \Omega\left(\frac{n^2}{(\log \log l)^2}\right) \geq \Omega\left(\frac{n^2}{(\log \log n)^2}\right)$. \square

3.5 Friends and Fiends

In some online networks, as well as in real life, one can be connected not only to one's friends but also to one's fiends (e.g. Epinions, Slashdot). This phenomenon is especially well known by parents of teenage kids, trying to influence their offspring. We model such a network by allowing not only positive but also negative influence between members. Informally, one can think of a negative link between u and v as u 's tendency to adopt the opinion opposite to that of v .

The proof given in [91] for the upper bound, as well as the lower bound construction shown in the last section, can be applied to this model. The definition of a "failed edge" has to be updated to apply also for negative edges. This is done in a straightforward manner by using the same definition for successful and failed edges in the case of positive ties and by using the opposite definition in case of negative ties. Namely, a negative outgoing edge $\langle v, u \rangle$ fails on time step t if u adopts v 's opinion on time step $t + 1$, and succeeds otherwise.

We get the following results.

Lemma 3.21. *There exists a family of n -node unweighted synchronous influence networks with stabilization time $\Omega(n^2 / \log^2 n)$.*

Lemma 3.22. *An n -node unweighted influence network with positive and negative friendship ties stabilizes in $O(n^2)$ time steps.*

3.6 Weighted Influence Network

In a social network it seldom happens that all ties have the exact same interpretation. Considering, for instance, different acquaintance ties between people, one may observe that usually people listen to their best friends more than to their colleagues. We model the influence between a pair of nodes by assigning a weight to the corresponding edge. It is then assumed that a node changes its opinion if the *weighted majority* of

its neighbors have a different opinion. We are interested in the influence of adding weights to our graph on the stabilization time. We start by proving the following lemma.

Lemma 3.23. *An n -node weighted influence network G with assigned edge weights $w \in \mathbb{N}$ stabilizes in $\min\{2\omega(G), 2^n\}$ time steps.*

Proof. Note that there is a bijective relation between weighted graphs and multigraphs. A weighted graph can be modeled as a multigraph by replacing each edge e of weight $\omega(e) = k$ by k edges of weight 1 each. Conversely, each multigraph can be modeled as a weighted graph with weights $\omega(e) \in \mathbb{N}$ by substituting k multiedges by a single edge of weight k . The transformation does not influence the behavior of the nodes as in both situations the weight of the influence is not changed. For the multigraph we can apply Theorem 3.5 to conclude that the multigraph stabilizes in $2|E|$ time steps. As the number of edges in the multigraph corresponds to the weight $\omega(G)$ of the weighted graph, we conclude that a weighted IN stabilizes in $2\omega(G)$ time. Moreover, this process is deterministic, and the execution enters a cycle once some global state repeats itself. Consequently, since the IN has 2^n global states, it must stabilize in at most 2^n time step. \square

The stabilization time of an IN can not be prolonged arbitrarily by just setting the edge weights higher. A path network, for example, will stabilize in $O(n)$ rounds no matter how the edge weights are chosen. It is an intriguing question if the weights do indeed have an influence on the stabilization time or if there is another mechanism which may prevent INs from having a higher stabilization time than $O(n^2)$. As we show in the next paragraphs, edge weights can significantly increase the stabilization time of INs. We do this by presenting a family of graphs with stabilization time $2^{\Omega(n)}$.

Lemma 3.24. *There is a family of n -node weighted influence networks with stabilization time $2^{\Omega(n)}$.*

To provoke as many changes as possible we build a graph consisting of 3 different component types: Two different colored paths of length l and

several levels of a structure to which we refer as *transistor lines*. The transistor lines consist of 2 different colored lines of k transistors. The main idea is that the paths, with a suitable initial coloring, get “discharged” by a lengthy process during which they change their colors node by node as often as possible, and once a path is completely discharged, it gets recharged (i.e., reset to the original color pattern) by the transistor above it. In turn, each transistor in the transistor lines recharges the levels below it. So each level adds another multiplicative factor to the stabilization time. (Let us remark that we have programmed the construction and simulated the influence propagation process on this graph; the interested reader may find a program and a video tracing the simulation at <http://www.disco.ethz.ch/members/barkelle/FUN.zip>) Let us now take a closer look how the two paths work.

At round 0 all the nodes in path P^1 are blue and all the nodes in path P^2 are red. The first nodes of the paths are denoted by F . When they change their color they start a cascade of changes through the path. To achieve this, the weights of the edges between the path nodes are decreasing from the first to the last node. This ensures that the change of the first node is cascaded through the path without any influence going the opposite direction. The summed up influence to change the color of the first node has therefore also to be higher than the weight of the edge between the first and the second node. The paths are illustrated in Figure 3.12.

Definition 3.25. We define our path graph $P = (V_P, E_P)$ as an undirected weighted graph, where $V_P = \{p_1, \dots, p_k\}$ and

$$E_P = \{(p_i, p_{i+1}, 2k + l - i) \mid i = 0, 1, \dots, k - 2\}.$$

The levels above the paths consist of two transistor lines. At time 0, line L^1 is blue and L^2 is red. Each transistor line is composed of k transistors. The basic function of the transistor is to change the color in the level below it in a controlled manner. A transistor (see Figure 3.13) consists of three nodes: A switch node (Sw), a collector node (Co) and an emitter node (Em). The idea behind this is to control the color of the transistor by the switch node. This is done by using the switch node to change the color of the collector node which in return changes the color




Figure 3.12: The path nodes are connected with decreasing weights from the first node (F) to the last (L) in the path. This induces a cascade of color changes through the path once the first node changes its opinion. The edge between the last and second-to-last node has a weight larger than $k \cdot 2$ where k is the number of transistors in a transistor line, in order to prevent influence from evolving along the wrong direction. The cascade of changes is triggered by changing the color of one external node that is connected to F .




Figure 3.13: A transistor consisting of the following three nodes: switch (Sw), collector (Co) and an emitter (Em). Its edge weights satisfy the following equations:

- (1) $2x > \sum_{(u, Em) \in E} \omega(u, Em)$,
- (2) $x < y < x + 3$.




Figure 3.14: These four nodes will never change their color, as they share an edge with a higher weight than all the other adjacent edges combined. With k transistors per line, this weight is set to $k \cdot y + 1$.




Figure 3.15: A level of transistor lines, consisting of k blue and k red transistors. Each switch is wired to the emitter of the transistor in front of it, so that the transistors get activated in the order from left to right.

of the emitter node. To do so, the switch node shares an edge with the collector node of weight 3 and the collector node is balanced in the way that it will only change its color, if the summed up influence from the level above is the opposite color and the switch node changes to this color. The collector node shares an edge with the emitter node which is heavier than all the other edges adjacent to the emitter node combined. This makes sure that the emitter node changes its color exactly one round after the collector node changed its color, no matter what the colors of the other neighboring nodes are. The order in which the transistors are activated is level by level, and in each level - transistor by transistor. To make sure that a transistor is only activated when the transistor in front of it already finished, we add an edge of weight 2 between the emitter node and the switch node of the next transistor in the transistor line (the switch node of the first transistor in each line is connected to the last emitter of the opposite colored transistor line); see Figure 3.15.

Above the levels we have 4 special nodes, $S_1 - S_4$. These 4 nodes

consist of two pairs of nodes, one (S_1, S_2) blue and the other (S_3, S_4) red. Each pair is connected by an edge which is heavier than all the other edges adjacent to these nodes combined (see Figure 3.14). This ensures that these nodes never change their color and can be used as a stable color reservoir.

Definition 3.26. A transistor $T = (V_T, E_T)$ is an undirected weighted graph with

$$\begin{aligned} V_T &= \{Sw, Co, Em\} \\ E_T &= \{(Sw, Co, 3), (Co, Em, x)\}, \end{aligned}$$

where x depends on the level i of the transistor and the length l of the path. Furthermore,

$$\begin{aligned} x_0 &= k \cdot (k \cdot 2 + l) + 2k + 3 \\ x_i &= k \cdot (x_{i-1} + 2k + 3) + 3. \end{aligned}$$

Definition 3.27. A transistor line $L^j = (V_{L^j}, E_{L^j})$ is an undirected weighted graph consisting of k transistors, where

$$\begin{aligned} V_{L^j} &= \bigcup_{i=0}^{k-1} V_{T_i} \\ E_{L^j} &= \bigcup_{i=0}^{k-1} E_{T_i} \cup \{(Em_i, Sw_{i+1}, 2) \mid i = 0, 1, \dots, k-1\}. \end{aligned}$$

Definition 3.28. A level $L = (V_L, E_L)$ is an undirected weighted graph consisting of two transistor lines L^1, L^2 , where $V_L = V_{L^1} \cup V_{L^2}$ and

$$\begin{aligned} E_L &= E_{L^1} \cup E_{L^2} \cup \{(u, v) \mid u = Sw \in T_0 \in L^1 \wedge v = Em \in T_{k-1} \in L^2\} \\ &\quad \cup \{(u, v) \mid u = Sw \in T_0 \in L^2 \wedge v = Em \in T_{k-1} \in L^1\}. \end{aligned}$$

We now show how the structures and different levels are wired. We start with the two paths and the first transistor lines. We want the blue path P^1 to turn alternately red and blue. As the blue transistors have the potential to turn red we connect the emitter of the first blue transistor with the first node of the blue path with weight $w = 2 \cdot k + l$. Similarly, we connect the first red transistor with the first node of the red path. In order to turn them back to their original color we connect the emitter of the second red transistor to the first node of P^1 and the emitter of the second blue transistor to the first node of P^2 . We continue this until the first node of P^2 is connected with all the emitters of the even transistors in L^1 and all the emitters of the odd transistors in L^2 . To inhibit the second transistor from changing P^1 back before the first cascade finished we connect the last node in P^2 with the switch of the second transistor in L^2 with a weight $w = 2$. So the switch node can only switch if the transistor in front of it and the last node of the opposite colored path did switch. As P^1 and P^2 have the same length and start at the same time they will also finish at the same time the cascade and will influence the second transistors to switch. All these edges are added for the first node F in P^2 and the last node L in P^2 respectively. Note that with k being odd, there is always a summed up influence on the first node of the path with value $2k + l$. The edges added between the paths and the first levels can be summarized by:

$$\begin{aligned}
 E_{PL_1} = & \{(u, v, k') \mid u = F \in P^1 \wedge v \in \{Em_j \in L_1^1 \wedge j \text{ even}\} \cup \{Em_j \in L_1^2 \wedge j \text{ odd}\}\} \\
 & \cup \{(u, v, k') \mid u = F \in P^2 \wedge v \in \{Em_j \in L_1^2 \wedge j \text{ even}\} \cup \{Em_j \in L_1^1 \wedge j \text{ odd}\}\} \\
 & \cup \{(u, v, 2) \mid u = L \in P^1 \wedge v \in \{Sw_j \in L_1^2 \wedge j \text{ even}\} \cup \{Sw_j \in L_1^1 \wedge j \text{ odd}\}\} \\
 & \cup \{(u, v, 2) \mid u = L \in P^2 \wedge v \in \{Sw_j \in L_1^1 \wedge j \text{ even}\} \cup \{Sw_j \in L_1^2 \wedge j \text{ odd}\}\},
 \end{aligned}$$

where $k' = 2k + l$.

The different levels are wired similarly to the first level and the path. The first node of the path corresponds to the collector nodes of the transistors one level below and the last node of the path corresponds to the emitter node of the last transistor. Formally we add the following edges between level L_i and level L_{i+1} .

$$\begin{aligned}
E_{LL_i} = & \{(u, v, x_i) \mid u = Co_j \in L_i^1 \wedge v \in \{Em_m \in L_{i+1}^2 \wedge j, m \text{ even}\}\} \\
& \cup \{(u, v, x_i) \mid u = Co_j \in L_i^1 \wedge v \in \{Em_m \in L_{i+1}^2 \wedge j, m \text{ odd}\}\} \\
& \cup \{(u, v, x_i) \mid u = Co_j \in L_i^2 \wedge v \in \{Em_m \in L_{i+1}^1 \wedge j, m \text{ even}\}\} \\
& \cup \{(u, v, x_i) \mid u = Co_j \in L_i^2 \wedge v \in \{Em_m \in L_{i+1}^1 \wedge j, m \text{ odd}\}\} \\
& \cup \{(u, v, 2) \mid u = Sw_j \in L_{i+1}^1 \wedge j \text{ even} \wedge v = Em_{k-1} \in L_{i-1}^1\} \\
& \cup \{(u, v, 2) \mid u = Sw_j \in L_{i+1}^1 \wedge j \text{ odd} \wedge v = Em_{k-1} \in L_{i-1}^2\} \\
& \cup \{(u, v, 2) \mid u = Sw_j \in L_{i+1}^2 \wedge j \text{ even} \wedge v = Em_{k-1} \in L_{i-1}^2\} \\
& \cup \{(u, v, 2) \mid u = Sw_j \in L_{i+1}^2 \wedge j \text{ odd} \wedge v = Em_{k-1} \in L_{i-1}^1\}.
\end{aligned}$$

The last level L_r is wired to the special nodes in the following way:

$$\begin{aligned}
E_{SL_r} = & \{(u, v, x_r + 2k + 2) \mid u = S_1 \wedge v \in \{Co_j \in L_r^2\}\} \\
& \cup \{(u, v, x_r + 2k + 2) \mid u = S_3 \wedge v \in \{Co_j \in L_r^1\}\}.
\end{aligned}$$

The complete weighed IN is a union of all these structures and can be seen with $k = 3$ and $l = 3$ and 5 levels in Figure 3.16.

Definition 3.29. Our worst case IN is an undirected weighted graph consisting of 2 paths, r levels and the 4 special nodes. Formally it is defined as $IN = (V_{IN}, E_{IN})$, where

$$\begin{aligned}
V_{IN} &= V_{P1} \cup V_{P2} \cup \bigcup_{i=1}^r V_{L_i} \cup \{S_1, S_2, S_3, S_4\} \\
E_{IN} &= E_{P1} \cup E_{P2} \cup E_{PL_1} \cup \bigcup_{i=1}^{r-1} E_{LL_i} \cup E_{SL_r}
\end{aligned}$$




Figure 3.16: This is an IN with stabilization time $l \cdot k^{(n-l)/k*6}$ with $k = 3$, $l = 5$ and $r = 5$.

Stabilization Time

We analyze the stabilization time of the presented graph. Each time the path is activated it takes l rounds to complete. As the transistors from the first level are connected to the last nodes of the two paths they only change when all the nodes in the path changed their color. As a transistor needs 3 rounds to change (switch \rightarrow collector \rightarrow emitter) and nothing of this happens in parallel, the first level takes $k \cdot (3 + l)$ rounds. Each additional level adds a factor of k to the stabilization time which leads to the following recursive function for the running time: $T_i = k^i \cdot (3 + l) + T_{i-1}$. Solving the recurrence formula gives

$$T_i = \frac{(l + 3)(k^{i+1} - 1)}{k - 1} \in O(k^i).$$

The running time grows exponentially in the number of levels. Each level consists of $2k$ transistors and each transistor consists of 3 nodes. The constant nodes consist of $2 \cdot 2$ nodes. The IN consists of n nodes, therefore we can build an IN with $i = \frac{n-l-4}{6 \cdot k}$ levels. Choosing l to be constant and $k = 3$, we achieve a stabilization time of $\Omega(3^{n/18}) = \Omega(2^{0.088n})$.

3.7 Asymmetric Influence Network

In real life, ties between people do not necessarily have the same weight for both adjacent nodes. Although friendships are often symmetrically perceived concerning how strong they are, there are a lot of examples where this is not the case. One example is the student advisor relation. Usually the advisor's opinion has a larger influence on the student than vice versa, hence the edge between advisor and student has a smaller weight for the advisor than for the student. This is even more extreme in the case of celebrities: A famous artist may influence people whom she does not even know, who in return do not influence her at all. We extend the model to allow asymmetric weights. Note that the weight can also be 0 on one side, which is then equivalent to a directed edge. Interestingly, in this new model the "stable states" are not that simple anymore, as the cycle length can be larger than 2. We are interested in the cycle length which can be achieved. An easy lower bound on it is n . One can think

of a circle with edges directed in one direction. Initially, one node is red and all the others are blue. This red “token” cycles through the circle with cycle length n . We are interested in how big the cycle length can get in an IN with asymmetric weights. As asymmetric INs are deterministic too, we have an upper bound of 2^n . We show a family of graphs with a cycle length of $2^{\Omega(n)}$

Lemma 3.30. *There are families of n -node influence networks with a cycle length $2^{\Omega(n)}$.*

We use the same IN as described in Section 3.6 as a basis for our construction but substitute most of the symmetric edges by directed edges. The main idea is to have the same process as the IN in Section 3.6, except that in the round where the symmetric IN would stabilize, our IN gets restarted by changing the colors of the special 4 nodes $S_1 - S_4$. This leads to a cycle length for our asymmetric IN that’s twice as long, as the stabilization time of the previous IN. In order to do so, we add directed edges from the last emitter of each transistor line and each path to the nodes $S_1 - S_4$ with a weight x that sums up to a weight higher than the edge weight $w(S_1, S_2) = w(S_3, S_4)$ but so that each subset of these weights is smaller than w . This will change the special nodes exactly when all the levels have switched. This is achieved by assigning the edges $w(S_1, S_2) = w(S_3, S_4) = 3(r + 1) - 1$, where r is the number of levels. Note that as directed edges can be used, we do not need an exponential growth of the weights anymore which makes the graph simpler.

To build our *AIN* we change the edges in the IN from Sect. 3.6 in the following way: The edges in the path graph are directed from the first node (F) to the last (L) and are assigned weight 1. The edges between the emitter from the first level and F are directed towards F and have weight 4. For each transistor the two edges of the switch node with weight 2 are now directed towards the switch node with weight 2. The edge between the switch node and the collector node is substituted by a directed edge to the collector node with weight 3. The edge between a collector and an emitter in the same transistor stays symmetrical but its weight is now 4. All the edges between collector and emitter nodes from different levels are substituted by directed edges from the emitter of the higher level to the collector on the lower level with weight 4. All the edges between the

special nodes and the collectors from level r are now directed towards the collector nodes and have weight 4. The edge between S_1 and S_2 (as well as between S_3 and S_4) is symmetric and has an assigned weight of $3(r+1) - 1$. Additionally we add the edges E_{AIN} . The graph with the added edges and new values is illustrated in Figure 3.17

Definition 3.31. The additional edges are formally described as:

$$\begin{aligned}
E_{AIN} = & E_{IN} \cup \{(u, v, 3) \mid u = L \in P^1 \wedge v = S_1\} \\
& \cup \{(u, v, 3) \mid u = L \in P^2 \wedge v = S_3\} \\
& \cup \{(u, v, 3) \mid u = L \in P^1 \wedge v = S_2\} \\
& \cup \{(u, v, 3) \mid u = L \in P^2 \wedge v = S_4\} \\
& \cup \bigcup_{i=1}^r \{(u, v, 3) \mid u = Em_k \in L_i^1 \wedge v = S_1\} \\
& \cup \bigcup_{i=1}^r \{(u, v, 3) \mid u = Em_k \in L_i^2 \wedge v = S_3\} \\
& \cup \bigcup_{i=1}^r \{(u, v, 3) \mid u = Em_k \in L_i^1 \wedge v = S_2\} \\
& \cup \bigcup_{i=1}^r \{(u, v, 3) \mid u = Em_k \in L_i^2 \wedge v = S_4\}.
\end{aligned}$$




Figure 3.17: This n -node asymmetric IN has a cycle length of $2^{\Omega(n)}$

4

Conclusion

In this work, we presented simplified versions of the social phenomenon of influence. On the one hand, we presented evidence that gender influences the choices in student-mentor networks and we showed that the glass ceiling effect is one of the consequences of this behavior. We presented a simple model of a network formation process with three simple characteristics: a smaller entry rate for women, preferential attachment, and homophily. We proved that already these three characteristics are sufficient to produce a glass ceiling in the formed network. We are aware that this is a very abstract representation of a growing network and that many more factors come into play when generating a student-mentor network. A complex social phenomenon like the glass ceiling can never be fully explained with such a simplified model but we are intrigued by the fact, that a model with only these three characteristics already leads to a glass ceiling, a prominent symptom of the gender stereotypes present

in our society. We hope that the insights gained during our research helps to further investigate this phenomenon and its reasons in order to find appropriate counter measurements. We do believe that the work at hand presents a good basis to compare different counter measurements. Our research suggests, for example, that putting effort into increasing the number of women starting a career in computer science is more fruitful than defining a quota for women in leading positions. Of course, these counter measurements are not independent of each other. We believe that having role models in leading positions in computer science does increase the sense of belonging to computer science for young women. This, in turn, will increase the number of women starting a career in computer science and therefore work towards the goal of increasing the entry rate for women in the network.

On the other hand, we showed how long it can take until people converge to their opinion if they are constantly influenced by other people. We are aware that our model is not by itself suitable to explain every detail of real world phenonema, but instead we aim to give a elementary and above all mathematically tractable abstraction of reality. Surely, hardly anyone is only interested in having the same opinion as the majority of her friends, independently of the opinion itself. But the tendency to adapt to the opinion of friends or other influential persons cannot be denied. Parts of it happens during discussions with friends where new points of view are encountered and opinions readjusted and other parts can be explained with the desire to be similar to ones friends. The asymmetric social influence model is probably the best fit for society. Especially important leader figures have an asymmetric influence on people. So the bad news is that in the worst case it takes exponentially long for people to converge to their opinion. The good news is that, since people often have symmetric relationships and tend to group themselves with people with similar opinions, the presented worst case graph is very unlikely to appear in practice.

All in all, our models are simplified versions of influence between people but we hope that our findings help to better understand the consequences influence can have in society and that we contribute in providing a theoretical background for discussions about influence.

Curriculum Vitae

- April 22nd, 1982** Born in Wetzikon, Switzerland
- 1989-2002** Primary school in Hinwil ZH, high school (Matura Typus E) at Kantonsschule Zürcher Oberland, Wetzikon ZH, Switzerland
- 2002-2009** Studies in computer science, ETH Zürich, Switzerland
- 2005-2011** Studies in Didactics, ETH Zürich, Switzerland,
- 2009** M.Sc. in computer science, ETH Zürich, Switzerland
- 2010-2015** PhD student, research and teaching assistant, Distributed Computing Group, supervised by Professor Roger Wattenhofer, ETH Zürich, Switzerland
- 2011** Didaktischer Ausweis, ETH Zürich, Switzerland
- 2013** Exchange year, Weizman Institute of Science, Rehovot, Israel
- 2015** PhD Degree, Distributed Computing Group, ETH Zürich, Switzerland

Publications

The following list contains all publications that I have co-authored during my PhD studies. The authors are listed in alphabetical order.

- [1] Homophily and the Glass Ceiling Effect in Social Networks. Chen Avin, Barbara Keller, Zvi Lotker, Claire Mathieu, Yvonne-Anne Pigolet and David Peleg. *6th Innovations in Theoretical Computer Science (ITCS)*, January 2015
- [2] How even Tiny Influence can have a Big Impact!. Barbara Keller, David Peleg and Roger Wattenhofer. *7th International Conference on Fun with Algorithms (FUN)*, July 2014
- [3] Convergence in (Social) Influence Networks. Silvio Frischknecht, Barbara Keller and Roger Wattenhofer. *27th International Symposium on Distributed Computing (DISC)*, October 2013
- [4] On the Feasibility of Opportunistic Ad Hoc Music Sharing. Barbara Keller, Philippe von Bergen, Roger Wattenhofer and Samuel Welten. *Nokia Mobile Developer Challenge Workshop (MDC)*, June 2012

Bibliography

- [1] Adler, F.R., Gordon, D.M.: Information Collection and Spread by Networks of Patrolling Ants. *The American Naturalist* (1992)
- [2] Afek, Y., Alon, N., Barad, O., Hornstein, E., Barkai, N., Bar-Joseph, Z.: A Biological Solution to a Fundamental Distributed Computing Problem. *Science* (2011)
- [3] Alon, N., Spencer, J.H.: *The Probabilistic Method*. John Wiley & Sons (2004)
- [4] Asuncion, A.U., Goodrich, M.T.: Turning Privacy Leaks into Floods: Surreptitious Discovery of Social Network Friendships and Other Sensitive Binary Attribute Vectors. In: *Workshop on Privacy in the Electronic Society (WPES)*. (2010)
- [5] Avin, C., Lotker, Z., Pigolet, Y.A., Turkel, I.: Core-Periphery in Networks: An Axiomatic Approach. *arXiv:1411.2242* (2014)
- [6] Barabási, A.L., Albert, R.: Emergence of Scaling in Random Networks. *Science* **286**(5439) (1999) 509–512
- [7] Berger, E.: Dynamic Monopolies of Constant Size. *Journal of Combinatorial Theory Series B* (2008)

- [8] Bickson, D., Tock, Y., Zymnis, A., Boyd, S.P., Dolev, D.: Distributed Large Scale Network Utility Maximization. International Conference on Symposium on Information Theory (ISIT) (2009)
- [9] Bock, S.J., Taylor, L.J., Phillips, Z.E., Sun, W.: Women and Minorities in Computer Science Majors: Results on Barriers from Interviews and a Survey. *Women* **14**(1) (2013) 143–152
- [10] Bramoullé, Y., Currarini, S., Jackson, M.O., Pin, P., Rogers, B.W.: Homophily and Long-run Integration in Social Networks. *Journal of Economic Theory* **147**(5) (2012) 1754–1786
- [11] Brass, D.J.: Men’s and Women’s Networks: A Study of Interaction Patterns and Influence in an Organization. *Academy of Management Journal* **28**(2) (1985) 327–343
- [12] Brin, S., Page, L.: The Anatomy of a Large-Scale Hypertextual Web Search Engine. In: WWW. (1998)
- [13] Ceci, S.J., Williams, W.M.: Understanding Current Causes of Women’s Underrepresentation in Science. *Proceedings of the National Academy of Sciences* **108**(8) (2011) 3157–3162
- [14] Chen, W., Yuan, Y., Zhang, L.: Scalable Influence Maximization in Social Networks under the Linear Threshold Model. In: *Industrial Conference on Data Mining (ICDM)*. (2010) 88–97
- [15] Chung, F.R.K., Lu, L.: *Complex Graphs and Networks*. Volume 107 of CBMS Regional Conference Series in Mathematics. AMS Bookstore (2006)
- [16] Cotter, D.A., Hermsen, J.M., Ovadia, S., Vanneman, R.: The Glass Ceiling Effect. *Social Forces* **80**(2) (2001) 655–681
- [17] Davies-Netzley, S.A.: Women above the Glass Ceiling Perceptions on Corporate Mobility and Strategies for Success. *Gender & Society* **12**(3) (1998) 339–355

- [18] Ding, W.W., Murray, F., Stuart, T.E.: Gender Differences in Patenting in the Academic Life sciences. *Science* **313**(5787) (2006) 665–667
- [19] Doerr, B., Fouz, M., Friedrich, T.: Social Networks Spread Rumors in Sublogarithmic Time. In: *Symposium on Theory of Computing (STOC)*. (2011)
- [20] Easley, D., Kleinberg, J.: *Networks, Crowds, and Markets*. Cambridge Univ Press **6**(1) (2010) 6–1
- [21] Eder, D., Hallinan, M.T.: Sex Differences in Children’s Friendships. *American Sociological Review* (1978) 237–250
- [22] *Education at a Glance: Organisation for Economic Co-operation and Development*. OECD (2012)
- [23] Etzkowitz, H., Kemelgor, C., Neuschatz, M., Uzzi, B., Alonzo, J.: The Paradox of Critical Mass for Women in Science. *Science* **266** (1994) 51–54
- [24] Eyring, A., Stead, B.A.: Shattering the Glass Ceiling: Some Successful Corporate Practices. *Journal of Business Ethics* **17**(3) (1998) 245–251
- [25] Falk, E., Grizard, E.: *The Glass Ceiling Persists: The 3rd Annual APPC Report on Women Leaders in Communication Companies*. The Annenberg Public Policy Center, University of Pennsylvania. Retrieved March 4 (2003) 2005
- [26] Faulkner, W., Lie, M.: Gender in the Information Society Strategies of Inclusion. *Gender, Technology and Development* **11**(2) (2007) 157–177
- [27] *Federal Glass Ceiling Commission: Solid Investments: Making Full Use of the Nation’s Human Capital*. US Government, Department of Labor. Washington, DC (1995)

- [28] Fisher, A., Margolis, J., Miller, F.: Undergraduate Women in Computer Science: Experience, Motivation and Culture. In: Proceedings of the Twenty-eighth SIGCSE Technical Symposium on Computer Science Education, ACM (1997) 106–110
- [29] Flocchini, P., Lodi, E., Luccio, F., Pagli, L., Santoro, N.: Dynamic Monopolies in Tori. *Discrete Applied Mathematics - Special Issue on International Workshop on Algorithms, Combinatorics, and Optimization in Interconnection Networks (IWACOIN)* (2004) 197–212
- [30] Floréen, P., Kaski, P., Polishchuk, V., Suomela, J.: Brief Announcement: Distributed Almost Stable Marriage. In: *Principles of Distributed Computing (PODC)*. (2010)
- [31] Frenkiel, N.: *The Up and Comers: Bryant Takes Aim at the Settlers-In*. Adweek, March (1984)
- [32] Fruchterman, T.M.J., Reingold, E.M.: *Graph Drawing by Force-Directed Placement. Software: Practice and Experience* (1991)
- [33] Gale, D., Shapley, L.S.: *College Admissions and the Stability of Marriage*. *The American Mathematical Monthly* (1962)
- [34] Gardner, M.: *Mathematical Games: The Fantastic Combinations of John Conway's new Solitaire Game "Life"*. *Scientific American* (1970)
- [35] Goles, E., Olivos, J.: *Periodic Behaviour of Generalized Threshold Functions*. *Discrete Mathematics* (1980)
- [36] Goles, E., Tchuente, M.: *Iterative Behaviour of Generalized Majority Functions*. *Mathematical Social Sciences* (1983)
- [37] Heider, F.: *Attitudes and Cognitive Organization*. *Journal of Psychology* (1946)

- [38] Hill, C., Corbett, C., St Rose, A.: Why So Few? Women in Science, Technology, Engineering, and Mathematics. ERIC (2010)
- [39] Hoefer, M.: Local Matching Dynamics in Social Networks. In: International Colloquium on Automata, Languages and Programming (ICALP). (2011)
- [40] Ibarra, H.: Homophily and Differential Returns: Sex Differences in Network Structure and Access in an Advertising Firm. *Administrative Science Quarterly* (1992) 422–447
- [41] Ibarra, H.: Paving an Alternative Route: Gender Differences in Managerial Networks. *Social Psychology Quarterly* (1997) 91–102
- [42] Kamada, T., Kawai, S.: An Algorithm for Drawing General Undirected Graphs. *Information Processing Letters* (1989)
- [43] Karp, R., Schindelhauer, C., Shenker, S., Vocking, B.: Randomized Rumor Spreading. In: *Foundations of Computer Science (FOCS)*. (2000)
- [44] Kaufmann, M., Wagner, D.: *Drawing Graphs: methods and models*. Springer (2001)
- [45] Kelman, H.: Compliance, Identification, and Internalization: Three Processes of Attitude Change. *Journal of Conflict Resolution* (1958)
- [46] Kempe, D., Kleinberg, J.M., Tardos, É.: Influential Nodes in a Diffusion Model for Social Networks. In: *International Colloquium on Automata, Languages and Programming (ICALP)*. (2005)
- [47] Kipnis, A., Patt-Shamir, B.: On the Complexity of Distributed Stable Matching with Small Messages. *Distributed Computing* (2010) 151–161
- [48] Kleinberg, J.M.: Hubs, Authorities, and Communities. *ACM Computing Surveys (CSUR)* (1999)
- [49] Kohonen, T.: *Self-Organization and Associative Memory*. Springer (1989)

- [50] Kostka, J., Oswald, Y.A., Wattenhofer, R.: Word of Mouth: Rumor Dissemination in Social Networks. In: 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO). (2008)
- [51] Kottis, A.P.: Women in Management: The “Glass Ceiling” and How to Break It. *Women in Management Review* **8**(4) (1993)
- [52] Lagesen, V.A.: The Strength of Numbers: Strategies to Include Women into Computer Science. *Social Studies of Science* **37**(1) (2007) 67–92
- [53] Lariviere, V., Ni, C., Gingras, Y., Cronin, B., Sugimoto, C.R.: Bibliometrics: Global Gender Disparities in Science. *Nature* **504** (2013) 211–213
- [54] Lazarsfeld, P.F., Merton, R.K., et al.: Friendship as a Social Process: A Substantive and Methodological Analysis. *Freedom and Control in Modern Society* **18**(1) (1954) 18–66
- [55] Leskovec, J., Singh, A., Kleinberg, J.: Patterns of Influence in a Recommendation Network. *Advances in Knowledge Discovery and Data Mining* (2006)
- [56] Leskovec, J., Huttenlocher, D.P., Kleinberg, J.M.: Signed Networks in Social Media. In: *Conference on Human Factors in Computing Systems (CHI)*. (2010) 1361–1370
- [57] Leslie, S.J., Cim pian, A., Meyer, M., Freeland, E.: Expectations of Brilliance Underlie Gender Distributions Across Academic Disciplines. *Science* **347**(6219) (2015) 262–265
- [58] Ley, M.: DBLP: Some Lessons Learned. *Proceedings of the VLDB Endowment* **2** (2009)
- [59] Ley, T.J., Hamilton, B.H.: The Gender Gap in NIH Grant Applications. *Science* **322**(5907) (2008) 1472–1474

- [60] Longo, P., Straehley, C.J.: Whack! I've Hit the Glass Ceiling! Women's Efforts to Gain Status in Surgery. *Gender Medicine* **5**(1) (2008) 88–100
- [61] Maccoby, E.E.: *The Two Sexes: Growing Up Apart, Coming Together*. Harvard University Press (1998)
- [62] Manouchehr, Z.: On Dynamic Monopolies of Graphs with General Thresholds. *Discrete Mathematics* (2012)
- [63] Martell, R.F., Lane, D.M., Emrich, C.: *Male-female Differences: A Computer Simulation*. (1996)
- [64] McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology* (2001) 415–444
- [65] Mislove, A., Marcon, M., Gummadi, K.P., Druschel, P., Bhattacharjee, B.: Measurement and Analysis of Online Social Networks. In: 7th ACM SIGCOMM Conference on Internet Measurement. (2007) 29–42
- [66] Mitzenmacher, M., Upfal, E.: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press (2005)
- [67] Moss-Racusin, C.A., Dovidio, J.F., Brescoll, V.L., Graham, M.J., Handelsman, J.: Science Faculty's Subtle Gender Biases Favor Male Students. *Proceedings of the National Academy of Sciences* **109**(41) (2012) 16474–16479
- [68] Nagel, K., Schreckenberg, M.: A Cellular Automaton Model for Freeway Traffic. *Journal de Physique I* (1992)
- [69] Neumann, J.V.: *Theory of Self-Reproducing Automata*. University of Illinois Press (1966)
- [70] Othman, M., Latih, R.: Women in Computer Science: No Shortage Here! *Commun. ACM* **49**(3) (March 2006) 111–114

- [71] Pearl, J.: Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach. In: Second National Conference on Artificial Intelligence (AAAI). (1982)
- [72] Peleg, D.: Local Majority Voting, Small Coalitions and Controlling Monopolies in Graphs: A Review. In: In Proc. of 3rd Colloquium on Structural Information and Communication Complexity. (1996) 152–169
- [73] Peleg, D.: Size Bounds for Dynamic Monopolies. *Discrete Applied Mathematics* (1998)
- [74] Pell, A.N.: Fixing the Leaky Pipeline: Women Scientists in Academia. *Journal of Animal Science* **74**(11) (1996) 2843–2848
- [75] Poljak, S., Sûra, M.: On Periodical Behaviour in Societies with Symmetric Influences. *Combinatorica* (1983) 119–121
- [76] Roberts, E.S., Kassianidou, M., Irani, L.: Encouraging Women in Computer Science. *SIGCSE Bull.* (2002) 84–88
- [77] Rolls, E.T., Treves, A.: *Neural Networks and Brain Function*. Oxford University Press, USA (1998)
- [78] Sassler, S., Glass, J., Levitte, Y., Michelmore, K.: The Missing Women in STEM? Accounting for Gender Differences in Entrance into STEM Occupations. In: Annual meeting of the Population Association of America Presentation. (2011)
- [79] Sauerwald, T., Stauffer, A.: Rumor Spreading and Vertex Expansion on Regular Graphs. In: *Symposium on Discrete Algorithms (SODA)*. (2011) 462–475
- [80] Sauerwald, T., Sudholt, D.: A Self-Stabilizing Algorithm for Cut Problems in Synchronous Networks. *Theoretical Computer Science* (2010)
- [81] Sheltzer, J.M., Smith, J.C.: Elite Male Faculty in the Life Sciences Employ Fewer Women. *Proceedings of the National Academy of Sciences* (2014)

- [82] Shen, H.: Mind the Gender Gap. *Nature* (2013)
- [83] Smith-Lovin, L., McPherson, J.M.: You are Who You Know: A Network Approach to Gender. *Theory on Gender/Feminism on Theory* (1993) 223–51
- [84] Spertus, E.: Why Are There So Few Female Computer Scientists? Technical report, Cambridge, MA, USA (1991)
- [85] Stephan, P.E., Levin, S.G.: Leaving Careers in IT: Gender Differences in Retention. *The Journal of Technology Transfer* **30**(4) (2005) 383–396
- [86] Stross, R.: What has Driven Women out of Computer Science? *New York Times* **15** (2008)
- [87] Tuma, N.B., Hallinan, M.T.: The Effects of Sex, Race, and Achievement on Schoolchildren’s Friendships. *Social Forces* **57**(4) (1979) 1265–1285
- [88] Watts, D.J.: A Simple Model of Global Cascades on Random Networks. *Proceedings of the National Academy of Sciences* (2002)
- [89] West, J.D., Jacquet, J., King, M.M., Correll, S.J., Bergstrom, C.T.: The Role of Gender in Scholarly Authorship. *PloS One* **8**(7) (2013) e66212
- [90] Whitecraft, M.A., Williams, W.M.: Why Aren’t More Women in Computer Science? In *Making Software: What Really Works, and Why We Believe*. (2011)
- [91] Winkler, P.: Puzzled: Delightful Graph Theory. *Communications of the ACM* (2008) 104
- [92] Wolfram, S.: *A New Kind of Science*. Wolfram Media (2002)