

Wireless networks: lecture notes

R. Combes

January 29, 2016

Contents

1	Wireless networks: a primer	9
1.1	Wireless vs wired	9
1.2	Some 802.11 terminology	10
1.3	802.11 PHY	11
1.3.1	Spectrum and transmit power	11
1.3.2	Access techniques	12
1.3.3	Rate adaptation	12
1.4	802.11 MAC	13
1.4.1	Frame transmission	13
1.4.2	Resource allocation	14
1.4.3	Distributed coordination function (DCF)	15
1.5	Modelling of wireless networks	16
1.5.1	Signal propagation and Interference	16
1.5.2	Traffic models	16
1.6	References	17
I	Mathematical Tools	19
2	Introduction to Markov Chains	21
2.1	Markov chains: definition	21
2.1.1	Definition	21
2.1.2	Homogenous Markov chains	22
2.1.3	Transition Matrix	22
2.1.4	Matrix notation	23
2.1.5	Graph notation	23
2.2	Stationary distribution and ergodicity	24
2.2.1	Stationary Markov chains	24
2.2.2	Stationary distributions	24
2.2.3	Full balance conditions	24
2.2.4	Strong Markov property	25

2.2.5	Transience and recurrence	27
2.2.6	Irreducibility	28
2.2.7	Stationary distribution: existence and uniqueness	29
2.2.8	Ergodicity	32
2.2.9	Aperiodicity	33
2.2.10	Convergence to the stationary distribution	34
2.3	Reversibility	37
2.3.1	Reversible Markov chains	37
2.3.2	Detailed balance conditions	38
2.3.3	Examples and counter-examples	39
2.3.4	Stationary distribution	40
2.3.5	A sufficient condition for reversibility	41
2.4	References	42
3	Markov chains: stability and mixing Time	43
3.1	Stability and the Foster-Liapunov criterion	43
3.1.1	Rationale: the ODE case	43
3.1.2	The Foster criterion	44
3.1.3	Martingales	45
3.1.4	Optional stopping theorem	46
3.1.5	Proof of Foster's criterion	47
3.1.6	An illustration in one dimension	49
3.2	Mixing time of Markov chains	49
3.2.1	Sampling from a stationary distribution	49
3.2.2	Example for two states	51
3.2.3	Exponential Mixing	52
3.2.4	Mixing Time	54
3.2.5	Standardizing distance	54
3.2.6	Coupling	55
3.3	References	56
4	Stochastic approximation	57
4.1	The basic stochastic approximation scheme	57
4.1.1	A first example	57
4.1.2	The associated o.d.e	58
4.1.3	Instances of stochastic approximation algorithms	59
4.1.4	Stochastic gradient algorithms	59
4.1.5	Distributed updates	60
4.1.6	Fictitious play	60
4.2	Convergence to the o.d.e limit	61
4.2.1	Assumptions	61

<i>CONTENTS</i>	5
4.2.2 The main theorem	62
4.3 Intermediate results	66
4.3.1 Ordinary differential equations	66
4.3.2 Martingales	67
II Analysis of wireless networks	69
5 The ALOHA protocol	71
5.1 Packet multiple access	71
5.2 ALOHA, i.i.d. Bernoulli model	72
5.3 Full buffer analysis	73
5.4 Stability Region of ALOHA	74
6 The CSMA protocol	79
6.1 The CSMA algorithm	79
6.1.1 CSMA principles	79
6.1.2 Variants	80
6.1.3 Physical and Virtual Sensing	81
6.1.4 Formal description	81
6.2 Performance of CSMA: Bianchi's model	82
6.2.1 The key assumption	82
6.2.2 Stationary distribution	83
6.2.3 Fixed point equation	85
6.2.4 Throughput	85
6.3 Engineering Insights	86
6.3.1 Large user regime	86
6.3.2 Window size	87
6.3.3 Small slot regime	88
6.4 Typical Parameter Values	88
7 Scheduling	89
7.1 Scheduling in constrained queuing systems	89
7.1.1 Constrained queuing systems	89
7.1.2 Stability region	90
7.2 Centralized scheduling: the Max-Weight algorithm	91
7.2.1 The Max-Weight algorithm	91
7.2.2 Throughput optimality	92
7.2.3 Computational complexity and message passing	94
7.3 Distributed scheduling: CSMA with Glauber Dynamics	94
7.3.1 The algorithm	94

7.3.2	Throughput Optimality	95
7.3.3	Iterative scheme	97
7.4	References	98
8	Capacity scaling of wireless networks	99
8.1	The model	99
8.1.1	Node locations	99
8.1.2	Transmission schedules	100
8.1.3	High probability events	100
8.2	Arbitrary networks on a circle	101
8.2.1	Capacity upper bound	101
8.2.2	Capacity lower bound	102
8.3	Random networks on a sphere	104
8.3.1	Connectivity	104
8.3.2	Source destination paths	104
8.3.3	Upper bound on the capacity	105
8.3.4	Lower bound on the capacity	105
8.3.5	Partitioning the domain into cells	106
8.3.6	Schedules	106
8.3.7	Each cell has at least one node whp	107
8.3.8	Routing	107
8.3.9	Capacity	108
8.4	Area of a disk on a sphere	108

Introduction

Foreword

This document is based on a course given to master's students on the topic of "Wireless Local Area and Ad Hoc Networks". The course was given in Supelec (France) in January 2015.

The target audience for this is engineering/math students with a basic understanding of the following topics

- elementary probability: conditioning, independence, convergence of random variables.
- modelling of wireless channels: distance-dependent path-loss, shadowing, and fast-fading.
- physical layer techniques: channel coding, multiple access schemes (TDMA, CDMA, FDMA etc.)

Chapter 1

Wireless networks: a primer

In this chapter we provide a short exposition of wireless networks standards and the underlying engineering problems.

1.1 Wireless vs wired

We will be mostly address two types of networks: WLANs and Ad-hoc networks. Both are a set of computers linked through a shared wireless channel with an important difference. In WLANs (802.11 -like), some nodes (known as access points) act as coordinators for association and resource allocation. In ad-hoc networks, there is no coordination entity, and all decisions (association, routing, resource allocation. are taken in a fully distributed manner. The most widespread type of WLANS are called WiFi networks, define by the IEEE 802.11 standard. IEEE 802.11 is standardized by the Wifi Alliance.

The main advantage of wireless networks (as opposed to wired) is the absence of wires (sic!) which allows easy management and mobility of nodes. Early users of wireless networks were the logistics industry, hospitals and the health care industry, and education (for instance colleges).

The main challenges of going from wired to wireless are the following:

- Scarcity of resources: the main resource is usually radio-frequency spectrum, which is both scarce and tightly regulated. WLANS use unregulated spectrum, so there is little hope in obtaining more spectrum. The only way to improve network performance is to improve the “signal processing” (modulation, coding, reception techniques, multi-access schemes, resource allocation, routing etc.)
- Unreliable medium: unlike copper cables and optical fibers, the wireless medium is highly unreliable, because of interference, fading and mobility. Therefore one needs to use coding, reception techniques and acknowledgements to ensure reliable data delivery.

- Security: the wireless medium is easily listened to by a malicious entity, and passive eavesdropping (sniffing) is undetectable. Therefore data must be encrypted and authentication must be used to avoid theft of data or identity.
- Topology changes: nodes are mobile, so that their physical location is changing over time. One needs to keep track of nodes to ensure connectivity of the network. Also, in wireless networks with high mobility, connectivity might be intermittent (ON/OFF).

Standards for WLANS are only concerned with PHY and MAC layer, so that cross-layer optimization is usually not standardized, although some proprietary implementations exist. In terms of quality of service, the PHY and MAC of standards such as 802.11 are *purely best-effort*, so that Quality of Service (QoS) never is taken into account, beyond successful delivery of data.

1.2 Some 802.11 terminology

Although our goal is not to cover the 802.11 standard in detail, we provide some terminology used in this standard here. There are several network entities, called as follows:

- "Stations": entities exchanging data (laptops, computers, smart phones)
- "Access points" (APs): gateways to a wired network
- "Wireless medium": the medium that carries data frames. It may be either Radio-Frequency or Infra-Red light
- "Distribution system/backbone": (only applies when there are several access points) the entity performing localization and routing. This entity is usually Ethernet-based (802.03). The connection between 2 APs might be wireless. In that case one talks about a wireless bridge.
- "Basic service Set" (BSS): a group of stations and (possibly) access points

There are mainly three types of networks, with the following denomination:

1. Infrastructure BSS (ad-hoc): a set of stations linked directly without an access point. Those networks are generally short lived.
2. Infrastructure BSS (classic): a set of stations associated to a common access point. Here the access point relays all communication. Any station attempting to enter such a network must go through a procedure called association. Association is initiated by a station and granted by the access point.
3. Extended service set (ESS): BSS's chained by a backbone network. An ESS is identified by a single SSID (Service Set Identifier), which acts as the "name" of the network.

The basic actions that a network can accomplish are called "services". Services granted by a network are the following:

- Distribution: move a frame from an AP to a station
- Integration: frame delivery to a non 802.11 network
- Association: register a station to an AP
- Reassociation: change the AP when Quality of Service (QoS) is poor
- Disassociation: termination of existing association
- Authentication: secure exchange of identity prior to sending data
- Deauthentication: termination of authentication
- Confidentiality: prevent eavesdropping through encryption
- MSDU delivery: delivery of data to the recipient
- Transmit power control: control of transmit power used by stations
- Dynamic frequency selection: detect and prevent interference (to other and from other systems e.g radar).

Nodes in a wireless network are, by definition mobile. When there are several access points available, a station should be associated to the access point which maximizes the received signal power. When a station exits the service area of an access point, a transition must occur to ensure that the station does not lose connectivity. There are three types of transitions:

- Absence of transition: a station stays in the service area of a given access point
- BSS transition: a station moves between 2 access points of the same ESS. This transition is in principle seamless and requires exchange of information between access points.
- ESS transition: a station moves between 2 access points of different ESSs. This transition is not seamless, and typically causes an interruption at the application level. For instance if the station is using a VoIP (Voice over IP) application, this transition should cause a call drop. Seamless transition requires special upper layer protocols not included in the 802.11 standard.

1.3 802.11 PHY

1.3.1 Spectrum and transmit power

Two main physical media can be used:

- Radio Frequency (RF) in unregulated frequency bands (at 2.4 GHz or 5 GHz). The 2.4 GHz frequency band suffers from microwave oven interference, and does not propagate under rain and long distances, but propagates through walls.

- Infra-Red (IR) light in the 300 GHz frequency band. The signal does not propagate through doors and walls. This is interesting from the point of view of security as this renders eavesdropping impossible. It should be noted that 802.11 using IR is fully standardized but has never been, to the best of our knowledge, implemented.

The typical transmit power in 802.11 networks is 20dBm, which corresponds to a range between 100 and 500 meters. Spectrum is the main limiting resource, and is controlled by controlled by the Federal Communications Commission (FCC) in the United States. The corresponding agency in France is the Agence Nationale des FRéquences (ANFR). The frequency bands used by 802.11 networks are called Industrial, Scientific and Medical (ISM) frequency bands. Those bands are unregulated, in the sense that anyone may transmit on those bands providing that the transmit power level is below a fixed threshold. As such, those bands are not interference free, and source of interference include cordless phones and microwave ovens.

1.3.2 Access techniques

Over the years, the 802.11 standards have evolved, and the access techniques used have roughly followed the evolution seen in cellular networks. As of now, the four possible access techniques have been used:

- Frequency Hopping: The frequency band is split into narrow frequency bands called subcarriers. Transmitters jump periodically from subcarrier to subcarrier in a deterministic pattern. Used in legacy 802.11.
- Direct Sequence Spread Spectrum: The frequency band is used by all transmitters in its entirety, and multiple access is done using code division, in a CDMA-like fashion. Used in legacy 802.11 and 802.11b
- Orthogonal Frequency Division Multiplexing (OFDM): The frequency band is split into subcarriers which act as parallel channels. All transmitters use the whole frequency band. Used in 802.11a/g.
- Multiple Input Multiple Output (MIMO)-OFDM: Transmitters and receivers are equipped with several antennas, which enables to use each antenna as an independent channel by using appropriate pre-coding at the transmitter. Used in 802.11n/ac.

1.3.3 Rate adaptation

A defining feature of wireless networks, as opposed to cellular networks is the rate adaptation mechanism. Rate adaptation is the function that lets a transmitter choose the modulation and coding scheme for each packet (which determine the data rate). The

achievable data rate with arbitrarily low probability of error is an increasing function of Signal To Noise Ratio (SNR) at the receiver, by Shannon's noisy channel coding theorem. Hence, ideally, the receiver should measure the SNR and send that information back to the transmitter before she transmits a packet. In 802.11 there is no such feedback, so that the transmitter must adjust the data rate based on the successes/failures of previous packets transmissions. A set of rates is defined by the BSS, and transmitters may change the data rate on a per-packet basis, including control packets. Recent standards use MIMO and OFDM, so that the number of available data rates is large. Designing a proper rate adaptation mechanism is critical for good performance.

1.4 802.11 MAC

1.4.1 Frame transmission

Modern computer networks are, for the most part, packet-switched networks, and 802.11 is no exception. Data is partitioned in small elementary units called "packets", which are then transmitted by the networks in an independent manner. In 802.11 packets are called "frames" and the typical frame size is a few thousand bytes. Since the physical medium is unreliable due to fading, interference and noise, frame transmission follows the following rules:

- Each unicast frame must be acknowledged
- If no Acknowledgement (ACK) for a frame is received, the frame is considered lost, a retry counter is incremented (1 counter per frame) and the frame will be retransmitted at a later time.
- Transmitters are responsible for retransmission
- Each frame updates the Network Allocation Vector (NAV), see below.
- When a higher level protocol attempts to send a packet larger than a threshold, the packet is fragmented into several frames which are sent separately. The threshold is called Maximal Transmission Unit (MTU).

When a frame is received, the receiver calculates a Cyclic Redundancy Check (CRC), and compares it to the value of the CRC contained in the header of the received frame. The CRC is a logical function of the received bits, ensuring that one may detect any error burst of length less than a fixed size. Typically, a CRC of length m can detect an error burst of length m or less. If the CRC is correct, an ACK frame is sent to the sender. Otherwise a NACK frame is sent and the frame is considered lost.

A frame contains the following elements: control data, NAV, receiver and sender MAC addresses, sequence number, payload (the actual data) and CRC.

Frames are not transmitted in a continuous manner, so that frames are not adjacent to avoid time synchronisation issues, and to allow sensing by other transmitters. The

spacing used follows a set of rules, and different spacing is used based on the frames sent type:

- (S) Short SIFS : RTS/CTS (see below) and ACK/NACK. Those frames have high priority.
- (P) PCF PIFS: used by the PCF (see below), any node using PCF that wants to seize the medium must wait for it to be idle for a time equal to PIFS.
- (D) DCF DIFS: used by the DCF (see below), any node using DCF that wants to seize the medium must wait for it to be idle for a time equal to PIFS.
- (E) Extended EIFS: used when there has been an error in the previous frame transmission.

It should be noted that the acronym IFS stands for Inter Frame Spacing.

1.4.2 Resource allocation

Contrary to cellular networks, in wireless local area networks and ad-hoc networks, the channel is shared between all nodes, and there is no central entity that takes care of the resource allocation. Resource allocation consists in determining, at any given time, which node may use which network resource (i.e subcarriers, antenna etc.). Proper resource allocation is necessary to prevent the adverse effects of excessive interference. In our setting, when a node transmits, it is unaware of whether or not other nodes are transmitting. The decision to transmit must be taken in a *distributed manner*. In fact, the channel allocation in Ethernet (IEEE 802.3) follows the same principle. In both cases, this design principle is chosen because of low complexity and cost.

Listen before talking

The simplest idea to perform distributed resource allocation is to implement “listen before you talk’. Each node senses the wireless channel before sending data. The most basic form of sensing is physical sensing: each node measures the received power level, and compares it to a threshold. If the received power is above the threshold, the channel is deemed busy, otherwise the channel is deemed idle. It should be noted that although physical sensing is a passive operation (the sensing node does not transmit anything), it consumes power, so that it should only be used when virtual sensing is not feasible (see below). When node A can sense transmissions of node B, we say that node A hears node B.

The network allocation vector (NAV) and virtual sensing

In order to reduce the need for physical sensing, another mechanism called NAV is used. When a node transmits a frame, the frame includes a numerical value called NAV, which

indicates the remaining amount of time she will use the channel. This is called virtual carrier sensing. Any node that hears a packet with a NAV strictly larger than 0 does not transmit nor senses physically the medium until a duration of NAV is elapsed. This limits the amount of energy consumed by physical sensing, which is critical for nodes that rely on batteries, for instance mobile phones.

The hidden node problem

Not only is physical sensing costly, but it is also subject to the so-called hidden node problem. Namely, consider nodes A, B and C placed on a line, with B between A and C. Nodes A and C both want to transmit to B. If the distance between them is large enough, A and C cannot hear each other, but B can hear both A and C. Hence cannot communicate but they interfere if they transmit simultaneously. The solution adopted in 802.11 is to use RTS /CTS (Request To Send / Clear To Send). If A wants to communicate with B, the following procedure is used:

- (a) A sends an RTS frame to B. The frame is heard by B, and B is silenced.
- (b) B sends back a CTS frame to A. The frame is heard by both A and C, C is silenced, and A starts transmitting.

In particular, when both sensing and RTS/CTS is used, then two nodes may collide (transmit simultaneously) during the RTS/CTS exchange. Since RTS/CTS frames are typically much smaller than a regular frame, the duration of a collision is much smaller than the duration needed for the successful transmission of a frame. This ensures that the medium is utilized efficiently.

1.4.3 Distributed coordination function (DCF)

Resource allocation in 802.11 is performed using the so-called DCF which is also called Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). Consider a node with a frame to transmit, and let n denote the retrial counter for this frame, i.e. the number of times that this frame has been transmitted without success. Let n_m be the maximal number of retransmissions.

1. The transmitter senses the medium, and waits until the medium is idle for a duration greater than DIFS. Both physical and virtual sensing are used.
2. The transmitter draws a random variable W uniformly distributed in $\{1, \dots, 2^{\min(n, n_m)}\}$ and sets $w = W$.
3. While $w > 0$: if the medium is sensed idle then w is decremented, otherwise w remains constant.
4. When $w = 0$, the packet is transmitted. And the transmitter waits for ACK/NACK.
5. If the transmission was successful, n is reset to 1. Otherwise n is incremented.

1.5 Modelling of wireless networks

1.5.1 Signal propagation and Interference

We quickly recall the standard models for the propagation of radio waves through a wireless channel. Consider a transmitter receiver pair located at distance r from each other. The transmitter transmits with power P , and the received power is modelled as $Pl(r)e^{\sigma Y}Z$ with:

- Path-loss: $l(r) = \min(Ar^{-\alpha}, 1)$, with $\alpha \geq 2$ and A a constant.
- Shadowing: Y a standard Gaussian variable and σ the shadowing standard deviation.
- Rayleigh Fading: Z and exponentially distributed random variable with parameter 1.

We recall that path-loss accounts for the distance-dependent loss due to propagation, shadowing accounts for absorption by obstacles such as walls, and fast-fading accounts for the fact that the signal typically travels through several uncorrelated paths where reflections on walls cause random phase changes.

Interference is described by two types of models:

- Protocol model: based on the transmitted power and the path loss exponent, one defines a transmission range for each node. When a receiver is within the transmission range of two nodes transmitting simultaneously, any transmission is considered unsuccessful. This model is a strong simplification of the physical reality, but allows to model a set of interfering nodes as a graph $G = (V, E)$. The vertices V are the nodes, and two nodes may transmit simultaneously if and only if they are not linked by an edge. As a consequence, a set of nodes may transmit simultaneously iff they form an independent set of G .
- Signal-To-Noise-plus-Interference (SINR) model: when a node transmits to another node, transmission is successful if and only if the SINR at the receiver is above a target value. This model is also called the physical model. Signal received from other transmitting nodes is treated as noise.

1.5.2 Traffic models

As in cellular networks, the main modelling problem comes from the fact that the state of buffers, the location of transmitting nodes and the location of users are time-varying. We will consider several types of models, corresponding to different time-scales:

- PHY-level model: there are 2 nodes, a transmitter and a receiver.
- MAC model, full-buffer model: there is one receiving node and several transmitting nodes. Transmitting nodes always have a packet to transmit at any given time

- MAC-level model, queuing model: there is one receiving node and several transmitting nodes. Transmitting nodes have a buffer where packets to be sent are stored. Packets arrive dynamically to the buffer of each node.
- Flow-level model: the number of nodes may vary across time. A node enters the network when it initiates a data flow (a series of packets to be transmitted). The node leaves the network when the last packet of her flow has been successfully delivered.

It should be noted that choosing the proper time scale is critical to ensure that the proposed models are both tractable and give a reasonable representation of the physical reality of the network.

1.6 References

The 802.11 standard is available at [15]. A more complete exposition of the standard and practical implementation of Wireless LANs is found in [11]. A comprehensive exposition of both wireless channel modelling can be found for instance in [30].

Part I

Mathematical Tools

Chapter 2

Introduction to Markov Chains

This chapter gives a short exposition of discrete Markov chains on countable state spaces. Markov chains are one of the fundamental building blocks for the performance evaluation and analysis of computer networks, including wireless networks.

2.1 Markov chains: definition

We consider $n \in \mathbb{Z}$. We consider a sequence $(X_n)_{n \in \mathbb{Z}}$ on a countable space \mathcal{X} . This sequence will be denoted $(X_n)_n$ unless it creates ambiguity. We identify the set of measures on \mathcal{X} with the set of positive vectors. For $x \in \mathcal{X}$ we denote by $\delta(x)$ the Dirac distribution at x , so that $\delta(x)_x = 1$ and $\delta(x)_{x'} = 0$ if $x' \neq x$. We denote by $\stackrel{(d)}{=}$ the equality in distribution between two random variables. We use the convention that $\inf \emptyset = \infty$.

2.1.1 Definition

Definition 2.1.1 (Markov property) $(X_n)_n$ is a Markov chain iff for all n and all $(x_0, \dots, x_n) \in \mathcal{X}^{n+1}$ one has:

$$\mathbb{P}[X_n = x_n | X_{n-1} = x_{n-1}, \dots, X_0 = x_0] = \mathbb{P}[X_n = x_n | X_{n-1} = x_{n-1}].$$

Defining feature of Markov chains is that the future and the past are independent conditionally on the present. Another way of phrasing this is to say that Markov chains are the process with order 1 memory: for all n , the distribution of X_n only depends on X_{n-1} .

2.1.2 Homogenous Markov chains

We will mainly consider homogenous Markov chains, where the transition probabilities are independent of n .

Definition 2.1.2 A Markov chain $(X_n)_n$ is homogenous iff for all $(x, x') \in \mathcal{X}^2$ and all n :

$$\mathbb{P}[X_n = x' | X_{n-1} = x] = \mathbb{P}[X_1 = x' | X_0 = x].$$

A homogeneous Markov chain can be interpreted as the stochastic version of a first order dynamical system in discrete time. In fact, $(X_n)_n$ is a Markov Chain iff there exists $f : \mathcal{X} \times [0, 1] \rightarrow \mathcal{X}$, and $(U_n)_n$ i.i.d uniformly distributed on $[0, 1]$ such that:

$$X_{n+1} = f(X_n, U_n).$$

2.1.3 Transition Matrix

Unless stated, we will consider homogeneous Markov chains throughout the chapter. We denote by μ the initial distribution with $\mu_x = \mathbb{P}[X_0 = x]$.

Definition 2.1.3 The transition matrix of a homogeneous Markov chain is $(P_{x,x'})_{x,x'}$ where:

$$P_{x,x'} = \mathbb{P}[X_1 = x' | X_0 = x].$$

Homogeneous Markov chains are a powerful tool in modelling, because, to define uniquely a Markov chain we only need to specify the transition matrix P and the initial distribution μ . Indeed, the probability of any event is expressed as a function of P and μ , since, by induction and the Markov property:

$$\begin{aligned} & \mathbb{P}[X_n = x_n, X_{n-1} = x_{n-1}, \dots, X_0 = x_0] \\ &= \mathbb{P}[X_n = x_n | X_{n-1} = x_{n-1}, \dots, X_0 = x_0] \mathbb{P}[X_{n-1} = x_{n-1}, \dots, X_0 = x_0] \\ &= P_{x_{n-1}, x_n} \mathbb{P}[X_{n-1} = x_{n-1}, \dots, X_0 = x_0] \\ &= \mu_{x_0} P_{x_0, x_1} \dots P_{x_{n-1}, x_n}. \end{aligned}$$

Based on this fact, we will call the (P, μ) -Markov chain the unique Markov chain with initial distribution μ and transition matrix P . Markov chains are ubiquitous in modelling of various systems, to cite but a few:

- The number of customers in a queuing system.
- The wireless channel
- The motion of dust particles on the surface of a liquid.
- Population dynamics in biology.
- The structure of DNA

- Modelling of textual documents
- Price of financial assets (e.g. stocks and interest rates)

We denote by $\mathbb{P}_x[\cdot] = \mathbb{P}[\cdot | X_0 = x]$, the probability of an event conditional to the fact that $X_0 = x$.

2.1.4 Matrix notation

Since \mathcal{X} is countable, we can identify μ with a line vector, and P with a square matrix. Both can have infinite dimension. Using the matrix notation, most calculations can be reduced to linear algebra, which is very practical. Let us summarize some facts:

Property 1 (i) If X_0 has distribution μ , X_1 has distribution μP (by the Markov property, see below),

(ii) If X_0 has distribution μ , X_n has distribution μP^n (by applying (i) n times),

(iii) Assume that $P^n \xrightarrow{n \rightarrow \infty} \bar{P}$, and X_0 has distribution μ . Then $(X_n)_n$ converges in distribution to $\mu \bar{P}$ (letting $n \rightarrow \infty$) in (ii).

(iv) Assume that $P^n \xrightarrow{n \rightarrow \infty} \bar{P}$, that \bar{P} has rank 1, with $\bar{\mu}$ the eigenvector associated to the unique non-null eigenvalue of \bar{P} . Then $(X_n)_n$ converges in distribution to $\bar{\mu}$ (applying (iii) for arbitrary μ).

(v) For all x, y , we have $\mathbb{P}[X_0 = x, X_n = y] = (P^n)_{xy}$ (applying (ii) with $\mu = \delta(x)$)

Regarding fact (iii) it should be noted that the sequence of matrices $(P^n)_n$ does not converge in general. For instance consider $P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, then $(P^n)_n$ does not converge since $P^2 = I$, which implies that $P^{2n+1} = P$ and $P^{2n} = I$ for all n .

Fact (iv) states that, in particular, if we can prove that the spectral radius of P is 1 and that all eigenvalues (but the largest) of P have absolute value strictly less than 1, then we can establish that $(X_n)_n$ converges in distribution (to the stationary distribution).

In particular, the distribution of X_1 is indeed μP , since:

$$\mathbb{P}[X_1 = x] = \sum_{x'} \mathbb{P}[X_1 = x, X_0 = x'] = \sum_{x'} \mu_{x'} P_{x',x} = (\mu P)_x. \quad (2.1)$$

2.1.5 Graph notation

It is often useful to represent the transition matrix of a Markov chain as a weighted directed graph $G = (V, E, W)$. Vertices are $V = \mathcal{X}$, and $(x, x') \in E$ iff $P_{x,x'} > 0$. The weight of edge (x, x') is $W_{x,x'} = P_{x,x'}$. We will denote by $G(P)$ the graph associated to P . The graph notation is useful because it allows to phrase properties of the transition matrix using graph theoretic notions, such as paths, cycles and so on.

We use the convention that the weight of a path $W(x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_m)$ is the product of weights of edges traversed so that it also represents the probability that the $(\delta(x_0), P)$ travels through this path:

$$W(x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_m) = P_{x_0, x_1} \dots P_{x_{m-1}, x_m} = \mathbb{P}[(X_0, \dots, X_m) = (x_0, \dots, x_m) | X_0 = x_0].$$

2.2 Stationary distribution and ergodicity

2.2.1 Stationary Markov chains

Simply said, a stochastic process is stationary iff its distribution does not change when time is shifted (by an arbitrary time shift).

Definition 2.2.1 A stochastic process $(X_n)_n$ is stationary iff for all $m \in \mathbb{N}$, we have

$$(X_n)_n \stackrel{(d)}{=} (X_{n+m})_n.$$

Namely for all n and all $(x_0, \dots, x_n) \in \mathcal{X}^{n+1}$:

$$\mathbb{P}[X_{n+m} = x_n, X_{n+m-1} = x_{n-1}, \dots, X_m = x_0] = \mathbb{P}[X_n = x_n, X_{n-1} = x_{n-1}, \dots, X_0 = x_0].$$

2.2.2 Stationary distributions

In general, when P and μ are arbitrary, the associated (μ, P) Markov chain is *not* stationary. For instance, consider: $\mathcal{X} = \{0, 1\}$, $P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ and $\mu = \delta(0)$. Clearly $(X_n)_n$ is not stationary since $\mathbb{P}[X_0 = 0] = 1 \neq \mathbb{P}[X_1 = 0] = 0$.

Definition 2.2.2 Consider μ a distribution on \mathcal{X} . μ is a stationary distribution for P iff the (μ, P) Markov chain is stationary.

2.2.3 Full balance conditions

Proposition 1 (Full balance condition) A distribution μ on \mathcal{X} is a stationary distribution for P iff:

$$\mu = \mu P.$$

i.e. for all x :

$$\mu_x = \sum_{x'} \mu_{x'} P_{x', x}.$$

Proof. If μ is a stationary distribution $X_1 \stackrel{(d)}{=} X_0$. From subsection 2.1.4 this implies $\mu = \mu P$. The other implication is left as an exercise. \square

Note that in general: 1) there might be no stationary distribution and 2) there might exist several stationary distributions.

Example of 1): $\mathcal{X} = \mathbb{N}$, $P_{x,x+1} = 1$ for all $x \in \mathbb{N}$. Assume that μ is a stationary distribution, then we must have $\mu_x = \mu_{x+1}$, so that μ is constant. Since μ is constant namely $\mu_x = \mu_0$ for all x . Then, if $\mu_0 = 0$ then $\mu = 0$ so that μ is not a probability distribution. If $\mu_0 > 0$ then $\sum_x \mu_x = \infty$ so that μ is not a probability distribution. Hence no stationary distribution exists.

Example of 2): $\mathcal{X} = \{0, 1\}$, P the identity matrix. Then any distribution on \mathcal{X} is stationary. We will later give conditions for uniqueness of the stationary distribution.

Also, a vector μ is said to be a stationary measure if it has positive elements and satisfies the full balance conditions. Note that we might have $\sum_x \mu_x = \infty$. Finally, if \mathcal{X} is finite and μ is a stationary measure, then $\frac{\mu}{\sum_x \mu_x}$ is a stationary probability.

2.2.4 Strong Markov property

Frequently, when manipulating Markov chains, one needs to consider events that occur at *random times*. For instance, given states x, x' , and initial distribution $\delta(x)$, we would like to know the (random) amount of time taken to reach state x' for the first time. The result that allows us to manipulate such events is the *strong* Markov property. We denote by \mathcal{F}_n the σ algebra generated by (X_0, \dots, X_n) . Simply said, \mathcal{F}_n is the family of all (random) events which can be expressed as a function of (X_0, \dots, X_n) . Namely a random variable Y is \mathcal{F}_n -measurable iff there exists a function g such that $Y = g(X_0, \dots, X_n)$. Also, \mathcal{F}_n can be seen as the information that an agent would have at time n if it observed the process $(X_{n'})_{n'}$ at times $n' \in \{0, \dots, n\}$.

Definition 2.2.3 A stopping time T is a random variable with values in $\mathbb{N} \cup \{\infty\}$ such that, for all $n \in \mathbb{N}$, we have:

$$\{T = n\} \in \mathcal{F}_n.$$

An alternative definition goes as follows: consider T a random variable with values in $\mathbb{N} \cup \{\infty\}$. T is a stopping time iff, for all n there exists a function g_n such that $\mathbf{1}\{T = n\} = g_n(X_0, \dots, X_n)$. Two examples of stopping times are:

- (deterministic stopping time) $T = n$ a.s. for a given $n \in \mathbb{N}$
- (first hitting time) $T = \inf\{n \geq 0 : X_n = x\}$ for a given $x \in \mathcal{X}$

Some counter-examples are (T is *not* a stopping time):

- (exit time) $T = \inf\{n \geq 0 : X_{n+1} \neq x\}$ for a given $x \in \mathcal{X}$
- (last hitting time) $T = \sup\{n \geq 0 : X_n = x\}$

Proposition 2 (Strong Markov property) Consider $(X_n)_n$ the (μ, P) -Markov chain and T a stopping time. Define the process $Y_n = (X_{T+n})_n$, with the convention that $Y_n = 0$ for all n if $T = +\infty$. Then we have that, for all $x \in \mathcal{X}$, conditionally to $\{T < \infty\} \cap \{X_T = x\}$, $(Y_n)_n$ is the $(\delta(x), P)$ Markov chain, and is independent of \mathcal{F}_T .

Proof. Consider $m \in \mathbb{N}$, and $(x_1, \dots, x_m) \in \mathcal{X}^m$. Consider $n \in \mathbb{N}$ fixed, and a sample path such that $T = n$ and $X_T = x$. Then $(Y_{T+1}, \dots, Y_{T+m}) = (x_1, \dots, x_m)$ iff $(X_{n+1}, \dots, X_{n+m}) = (x_1, \dots, x_m)$. Using the Markov property:

$$\mathbb{P}[(X_{n+1}, \dots, X_{n+m}) = (x_1, \dots, x_m) | X_n = x] = P_{x_{m-1}, x_m} \dots P_{x, x_1}.$$

Going back to process Y we have:

$$\mathbb{P}[(Y_{T+1}, \dots, Y_{T+m}) = (x_1, \dots, x_m) | X_n = x, T = n] = P_{x_{m-1}, x_m} \dots P_{x, x_1}.$$

The r.h.s. of the above equation does not depend on n , so that taking the union over $n \in \mathbb{N}$

$$\mathbb{P}[(Y_{T+1}, \dots, Y_{T+m}) = (x_1, \dots, x_m) | X_T = x, T < \infty] = P_{x_{m-1}, x_m} \dots P_{x, x_1}. \quad (2.2)$$

Two observations can be made to conclude the proof:

- Define $(Z_n)_n$ the $(\delta(x), P)$ Markov chain. The r.h.s. of (2.2) is equal to the probability that $(Z_1, \dots, Z_m) = (x_1, \dots, x_m)$. Hence conditionally to $\{X_T = x, T < \infty\}$, $Y \stackrel{(d)}{=} Z$. Therefore Y is the $(\delta(x), P)$ Markov chain.
- The r.h.s. of (2.2) does not depend on \mathcal{F}_T . So conditionally to $\{X_T = x, T < \infty\}$, Y is independent of \mathcal{F}_T .

□

Consider $x \in \mathcal{X}$ and define the hitting time $T_x = \inf\{n \geq 1 : X_n = x\}$. Assume that $T_x < \infty$ a.s. Then the strong Markov property simply states that, after time T_x the chain behaves as the $(\delta(x), P)$ Markov chain, and is independent of any event before T_x . This fact is used frequently to decompose a Markov chain into "excursions", which are random time epochs between two returns to a given state.

Corollary 2.2.4 Consider x such that $T_x < \infty - \mathbb{P}_x$ a.s. Define $T^0 = 0$ and for $k \geq 1$ the stopping times:

$$T_x^k = \begin{cases} \inf\{n \geq 1 : X_{n+T_x^{k-1}} = x\} & \text{if } T^{k-1} < \infty \\ \infty & \text{otherwise.} \end{cases}$$

Then random variables $(T_x^k)_{k \geq 1}$ are i.i.d (hence finite a.s.).

Proof. By definition $T_x^1 = T_x < \infty$ by assumption. We proceed by induction. Assume that we have proven for $K \geq 1$ that (T_x^1, \dots, T_x^K) is i.i.d. Then $T_x^K < \infty$ a.s. since $T_x^K \stackrel{(d)}{=} T_x$. By the strong Markov property, the process $(X_{n+T_x^k})_n$ is the $(\delta(x), P)$ Markov chain, and is independent of $\mathcal{F}_{T_x^k}$. Therefore $T_x^{K+1} \stackrel{(d)}{=} T_x$ and independent of T_x^1, \dots, T_x^K . By induction, we have proven that random variables $(T_x^k)_{k \geq 1}$ are i.i.d (hence finite a.s.), which concludes the proof. \square

2.2.5 Transience and recurrence

The strong Markov property and its corollary suggest that the hitting times T_x are particularly useful, and that the finiteness/non-finiteness T_x is important. It is natural to distinguish two types of states x depending on the finiteness/non-finiteness of T_x .

Definition 2.2.5 Consider $x \in \mathcal{X}$.

- (i) x is recurrent iff $\mathbb{P}_x[T_x = \infty] = 0$
- (ii) x is transient iff $\mathbb{P}_x[T_x = \infty] > 0$

Intuitively, recurrent states are (infinitely) often visited, while transient states are visited finitely many times. Define the number of visits to x on a sample path:

$$N_x = \sum_{n \geq 1} \mathbf{1}\{X_n = x\}$$

In fact the number of visits to transient states is geometrically distributed while the number of visits to a recurrent state is finite a.s. as shown by the following proposition.

Proposition 3 Define $p_x = \mathbb{P}_x[T_x < \infty]$.

- (i) If x is recurrent, $N_x = \infty$ \mathbb{P}_x - a.s.
- (ii) If x is transient, N_x is geometrically distributed with parameter p_x , under probability \mathbb{P}_x .

Proof. Define $T_x^1 = T_x$ and $T_x^2 = \inf\{n \geq 1 : X_{T_x+n} = x\}$. From the strong Markov property, T_x^2 is independent of T_x^1 and has the same distribution. Hence:

$$\mathbb{P}_x[N_x \geq 2] = \mathbb{P}_x[T_x^1 < \infty, T_x^2 < \infty] = (\mathbb{P}_x[T_x < \infty])^2 = p_x^2$$

By induction, for all $k \geq 1$:

$$\mathbb{P}[N_x \geq k] = p_x^k$$

We can conclude the proof by noting that:

- If x is recurrent, $p_x = 1$ so that $\mathbb{P}_x[N_x \geq k] = 1$ for all k , hence $\mathbb{P}_x[N_x < \infty] = 0$.
- If x is transient, $p_x < 1$ so that $\mathbb{P}[N_x = k] = (1 - p_x)p_x^k$, hence N_x has geometric distribution with parameter p_x .

In fact the expected number of visits to x can also be calculated directly using P . In subsection 2.1.4 we have shown that $\mathbb{P}_x[X_n = x] = (P^n)_{xx}$. Hence:

$$\mathbb{E}_x[N_x] = \sum_{n \geq 0} \mathbb{P}_x[X_n = x] = \sum_{n \geq 0} (P^n)_{xx}$$

Finally, we distinguish two types of recurrent states depending on the expected value of T_x . Intuitively, if T_x has infinite expectation, this means that x is visited infinitely many times, but the return time to x has large fluctuations. Indeed, from the Markov inequality, if T_x has a finite expectation, then for all N , $\mathbb{P}_x[T_x \geq N] \leq N^{-1} \mathbb{E}_x[T_x]$. We will later show that positive recurrence is one of the conditions used to establish existence of a stationary distribution.

Definition 2.2.6 Consider $x \in \mathcal{X}$ recurrent.

- (i) x is positive recurrent iff $\mathbb{E}_x[T_x] < \infty$
- (ii) x is null recurrent iff $\mathbb{E}_x[T_x] = \infty$

2.2.6 Irreducibility

We now define irreducibility, which intuitively means that given any two states x and x' , it is possible (with positive probability), to go from x to x' and back. Clearly, there exists Markov chains where it is not possible to go back and forth between arbitrary pairs of states (think for instance of a Markov chain whose transition matrix is the identity matrix).

Definition 2.2.7 Consider a transition matrix P and $G = G(P)$ the associated graph. x and x' communicate iff there exists a cycle containing both x and x' .

We write $x \leftrightarrow x'$ to denote that x and x' communicate. \leftrightarrow is an equivalence relation i.e it is symmetric, reflexive and transitive. This may be verified easily. The relationship \leftrightarrow can also be stated using matrix notation.

Proposition 4 $x \leftrightarrow y$ iff there exists $n_x, n_y > 0$ such that $(P^{n_x})_{xy}(P^{n_y})_{yx} > 0$.

Proof. There exists a path of positive probability going from x to y with length n_x iff $\mathbb{P}_x[X_{n_x} = y] = (P^{n_x})_{xy} > 0$. \square

Proposition 5 Consider C an equivalence class for \leftrightarrow . Then all states in C are either recurrent or transient.

Proof. We proceed by contradiction. If the property does not hold there exists x and x' in C , such that x is transient and x' is recurrent.

Since $x \leftrightarrow x'$ and $x' \leftrightarrow x$, there exists paths p (from x to x') and p' (from x' to x) of respective lengths m and m' and of respective weights w and w' . Consider $n = r + m + m'$ with $r \in \mathbb{N}$. The probability of event $X_n = x'$ is lower bounded by the probability that the chains starts at x at time 0, then follows path p to reach x' at time m , then follows a cycle to return to x' at time $m + r$, and finally travels along path p' back to x at time $m + r + m'$. We have that:

- $\mathbb{P}_x[X_m = x'] \geq w$,
- $\mathbb{P}_x[X_{m+r} = x' | X_m = x'] = \mathbb{P}_{x'}[X_r = x']$ (homogeneity)
- $\mathbb{P}_x[X_{m+r+m'} = x | X_{m+r} = x'] = \mathbb{P}_{x'}[X_{m'} = x] \geq w'$

Hence:

$$\mathbb{P}_x[X_n = x] \geq ww' \mathbb{P}_{x'}[X_r = x'].$$

Recall that for all x we have: $\mathbb{E}_x[N_x] = \sum_{n \geq 0} \mathbb{P}_x[X_n = x]$, so summing we have proven:

$$\mathbb{E}_x[N_x] \geq ww' \mathbb{E}_{x'}[N_{x'}]$$

If x is transient and x' is recurrent, then $\mathbb{E}_x[N_x] < \infty$ and $\mathbb{E}_{x'}[N_{x'}] = \infty$ which is a contradiction. Hence we have proven that all $x \in C$ are either recurrent or transient which completes the proof. □

Definition 2.2.8 Consider transition matrix P , P is irreducible iff there is a unique equivalence class for \leftrightarrow .

Proposition 6 Consider \mathcal{X} finite and P irreducible. Then all states $x \in \mathcal{X}$ are recurrent.

Proof. Left as an exercise (pigeon-hole principle) □

2.2.7 Stationary distribution: existence and uniqueness

Now we prove that any irreducible, positive recurrent transition matrix has a unique stationary probability. Interestingly enough, the proof for this result relies crucially on the decomposition of the Markov chain into excursions (and hence on the strong Markov property). The proof is constructive, so that we explicitly construct a stationary distribution.

Proposition 7 Consider a transition matrix P which is both irreducible and positive recurrent. Then P admits a unique stationary probability.

Proof.**Existence of a stationary distribution**

Consider $x \in \mathcal{X}$ fixed. Define $T_x = \min\{n \geq 1 : X_n = x\}$ the return time to x and

$$N_y = \sum_{n=1}^{T_x} \mathbf{1}\{X_n = y\},$$

the number of visits to y between two visits to x . Define the vector μ such that for all y : $\mu_y = \mathbb{E}_x[N_y]/\mathbb{E}_x[T_x]$. μ is well defined since $\mathbb{E}_x[T_x] < \infty$ and has positive elements. Furthermore:

$$\sum_y N_y = \sum_{n=1}^{T_x} \sum_y \mathbf{1}\{X_n = y\} = \sum_{n=1}^{T_x} 1 = T_x.$$

Taking expectations and dividing by $\mathbb{E}_x[T_x]$ this proves that μ is a probability distribution:

$$\sum_y \mu_y = 1.$$

Now let us prove that μ satisfies the full balance conditions. We have:

$$\mathbb{E}_x[N_y] = \sum_{n \geq 1} \mathbb{P}_x[X_n = y, n < T_x],$$

By the Markov property, we have:

$$\begin{aligned} \mathbb{P}_x[X_n = y, n < T_x] &= \sum_z \mathbb{P}_x[X_n = y, X_{n-1} = z, n-1 < T_x] \\ &= \sum_z P_{zy} \mathbb{P}_x[X_{n-1} = z, n-1 < T_x]. \end{aligned}$$

Replacing in the above equation we get:

$$\begin{aligned} \mathbb{E}_x[N_y] &= \sum_{n \geq 1} \sum_z P_{zy} \mathbb{P}_x[X_{n-1} = z, n-1 < T_x] \\ &= \sum_z P_{zy} \sum_{n \geq 1} \mathbb{P}_x[X_{n-1} = z, n-1 < T_x] \\ &= \sum_z P_{zy} \mathbb{E}_x[N_z]. \end{aligned}$$

Dividing by $\mathbb{E}_x[T_x]$ we obtain :

$$\mu_y = \sum_z P_{zy} \mathbb{E}_x[N_z].$$

since the above reasoning holds for all $y \in \mathcal{X}$, μ verifies the full balance equations, so that μ is indeed a stationary distribution.

Unicity of the stationary distribution

Assume that there exists another stationary measure λ which is non-null. We may assume that there exists x such that $\lambda_x = 1$ (otherwise we can multiply λ by a scalar factor).

By linearity λ is a stationary measure and verifies the balance conditions:

$$\lambda_y = \lambda_x P_{xy} + \sum_{z_1 \neq x} \lambda_{z_1} P_{z_1 y} = P_{xy} + \sum_{z_1 \neq x} \lambda_{z_1} P_{z_1 y}$$

Iterating the reasoning above we obtain:

$$\lambda_y = P_{xy} + \sum_{z_1 \neq x} P_{xz_1} P_{z_1 y} + \sum_{z_1, z_2 \neq x} \lambda_{z_2} P_{z_2 z_1} P_{z_1 y}.$$

Hence by induction:

$$\lambda_y \geq \sum_{n \geq 0} \sum_{z_1, \dots, z_n \neq x} P_{xz_n} P_{z_n z_{n-1}} \dots P_{z_1 y}$$

Now we notice that $P_{xz_n} P_{z_n z_{n-1}} \dots P_{z_1 y}$ is the probability that (conditional to $X_0 = y$), $(X_n)_n$ travels through the path $x \rightarrow z_n \rightarrow \dots \rightarrow z_1 \rightarrow y$, so that $X_{n+1} = y$ and $X_{n'} \neq x$ for all $n' \leq n$, which implies $T_x > n + 1$, hence :

$$\sum_{z_1, \dots, z_n \neq x} P_{xz_n} P_{z_n z_{n-1}} \dots P_{z_1 y} = \mathbb{P}_x[X_{n+1} = y, T_x \geq n + 1]$$

and replacing:

$$\lambda_y \geq \sum_{n \geq 0} \mathbb{P}_x[X_{n+1} = y, T_x \geq n + 1] = \mathbb{E}_x[N_y] = \mu_y \mathbb{E}[T_x].$$

Hence, both λ and $(\mu_y \mathbb{E}[T_x])_y$ are stationary measures, and so is $\gamma = \lambda - (\mu_y \mathbb{E}[T_x])_y$, by linearity. Also notice that $\gamma_x = \lambda_x - \mathbb{E}_x[N_y] = 1 - 1 = 0$. Consider $y \in \mathcal{X}$. Since γ verifies the balance conditions, for all n

$$0 = \gamma_x = \sum_z \gamma_z (P^n)_{zx} \geq \gamma_y (P^n)_{yx}$$

We can now use the irreducibility. Since there exists n such that $(P^n)_{yx} > 0$, so that from the above equation we have $\gamma_y = 0$. Since the reasoning above holds for all y we have $\gamma = 0$, so that $\lambda = (\mu_y \mathbb{E}[T_x])_y$. We have proven that all stationary measures are proportional to $(\mu_y \mathbb{E}[T_x])_y$, which implies that the stationary probability is unique. \square

2.2.8 Ergodicity

We may now establish the ergodic theorem for Markov chains, a fundamental result showing that (for irreducible positive recurrent chains), the frequency at which x is visited (over a large time horizon) is equal to its stationary probability μ_x almost surely. Namely define the empirical frequencies:

$$N_x(n) = \sum_{n'=0}^n \mathbf{1}\{X_{n'} = x\},$$

$$\hat{\mu}_x(n) = \frac{N_x(n)}{n}.$$

We state the ergodic theorem from Markov chains, which generalizes the law of large numbers for sums of i.i.d random variables.

Theorem 2.2.9 (*Ergodic theorem for Markov chains*) Consider $(X_n)_n$ an irreducible, positive recurrent Markov chain with unique stationary distribution μ . Consider μ' an arbitrary starting distribution. Then for all $x \in \mathcal{X}$ we have:

$$\hat{\mu}_x(n) \xrightarrow[n \rightarrow \infty]{} \mu_x \quad \mathbb{P}_{\mu'} - a.s.$$

Before proving the result, we prove an intermediate result on the return time to x starting from an arbitrary state y .

Proposition 8 Consider $(X_n)_n$ an irreducible, positive recurrent Markov chain. Define $T_x = \min\{n \geq 1 : X_n = x\}$. Then for all $y \in \mathcal{X}$ we have:

$$\mathbb{E}_y[T_x] < \infty.$$

Proof. It is noted that for $y = x$, the result holds simply because x is positive recurrent. By irreducibility there exists m and a path of positive probability $x = z_0 \rightarrow \dots \rightarrow z_m = y$ with $z_1 \neq x, \dots, z_m \neq x$. Define event $A = \{X_0 = z_0, \dots, X_m = z_m\}$. Define $T'_x = \min\{n \geq 1 : X_{n+m} = x\}$. If A occurs then $T_x = T'_x + m$. So:

$$\begin{aligned} \mathbb{E}_x[T_x] &\geq \mathbb{E}_x[T_x \mathbf{1}\{A\}] = \mathbb{E}_x[(T'_x + m) \mathbf{1}\{A\}] \\ &= \mathbb{E}_x[T'_x \mathbf{1}\{A\}] + m\mathbb{P}[A] = (\mathbb{E}_x[T'_x|A] + m)\mathbb{P}[A]. \end{aligned}$$

From the Markov property $\mathbb{E}_x[T'_x|A] = \mathbb{E}_y[T_x]$ since $X_m = y$ when A occurs. We have proven:

$$\mathbb{E}_x[T_x] \geq (\mathbb{E}_y[T_x] + m)\mathbb{P}[A].$$

Since $\mathbb{P}[A] > 0$ and $\mathbb{E}_x[T_x] < \infty$ the above equation proves that $\mathbb{E}_y[T_x] < \infty$ which concludes the proof. \square

We are now fully equipped to prove the ergodic theorem for Markov chains.

Proof. [of Theorem 2.2.9] First consider the chain starting at x , and define T_x^k the k -th passing time to x . We have:

$$\sum_{k=1}^{N_x(n)-1} T_x^k \leq n \leq \sum_{k=1}^{N_x(n)} T_x^k$$

and dividing by $N_x(n)$:

$$\frac{\sum_{k=1}^{N_x(n)-1} T_x^k}{N_x(n)} \leq \frac{1}{\hat{\mu}_x(n)} \leq \frac{\sum_{k=1}^{N_x(n)} T_x^k}{N_x(n)}$$

x is recurrent so that $N_x(n) \rightarrow \infty$ a.s, and $\frac{1}{K} \sum_{k=1}^K T_x^k \rightarrow \mathbb{E}_x[T_x]$ a.s. by the law of large numbers. Recall that (T_x^k) are i.i.d with finite expectation since x is positive recurrent. We have proven that $\hat{\mu}_x(n) \rightarrow \frac{1}{\mathbb{E}[T_x]} = \mu_x$ a.s

Now consider the case where the starting distribution is $\delta(y)$. Then:

$$T_x^1 + \sum_{k=2}^{N_x(n)-1} T_x^k \leq n \leq T_x^1 + \sum_{k=2}^{N_x(n)} T_x^k.$$

T_x^1 is the first hitting time of x starting from y , and we have $T_x^1 < \infty$ a.s since $\mathbb{E}_y[T_x^1]$ by proposition 8. Hence $T_x^1/N_x(n) \xrightarrow[n \rightarrow \infty]{} 0$ a.s. so that, by the same reasoning as above we have that $\hat{\mu}_x(n) \rightarrow \frac{1}{\mathbb{E}[T_x]} = \mu_x$ a.s. which concludes the proof. \square

2.2.9 Aperiodicity

As stated previously, the limiting distribution of $(X_n)_n$ when $n \rightarrow \infty$ need not exist, and a simple counter-example is $P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Also notice that P is irreducible and positive recurrent (the return time to any state is upper bounded by 2), and has stationary distribution $(1/2, 1/2)$. The main feature of P is that any sample path is periodic, and has a period of 2.

Therefore, it is natural to introduce aperiodicity, which is (roughly) a condition ensuring that the sample paths of the Markov chain of interest are not periodic.

Definition 2.2.10 Consider P a transition matrix.

(i) $x \in \mathcal{X}$ is aperiodic iff there exists $n_0 \geq 1$ such that $\mathbb{P}_x[X_n = x] = (P^n)_{xx} > 0$ for all $n \geq n_0$.

(ii) P is aperiodic iff x is aperiodic for all $x \in \mathcal{X}$.

In fact, for irreducible transition matrices, all states are aperiodic iff there exists a single aperiodic state, as stated by proposition 9. The third statement of proposition 9 can be rephrased as follows: any irreducible Markov chain for which $\mathbb{P}[X_n = X_{n+1}] > 0$ is necessarily aperiodic. As a consequence, any lazy Markov chain is aperiodic. We recall that transition matrix P is said to be lazy iff $\inf_x P_{xx} \geq 1/2$. Namely, a lazy Markov chain is a chain that does not move with probability at least $1/2$.

Proposition 9 Consider P an irreducible transition matrix.

- (i) P is aperiodic iff there exists $x \in \mathcal{X}$ aperiodic.
- (ii) Consider x aperiodic and $y \in \mathcal{X}$. Then there exists $n_{xy} > 0$ such that for all $n \geq n_{xy}$ we have $(P^n)_{xy} > 0$.
- (iii) If P has a non-null diagonal, then P is aperiodic.

Proof. (i) Consider x aperiodic and $y \in \mathcal{X}$. Since P is irreducible there exists n_x, n_y such that $(P^{n_x})_{xy} > 0$ and $(P^{n_y})_{yx} > 0$. Since x is aperiodic, there exists n_0 such that for all $n \geq n_0$ one has $(P^n)_{xx} > 0$. Consider $n \geq n_0$. By the Markov property:

$$\begin{aligned} (P^{n+n_x+n_y})_{yy} &= \mathbb{P}_y[X_{n+n_x+n_y} = y] \\ &\geq \mathbb{P}_y[X_{n_y} = x, X_{n_y+n} = x, X_{n+n_y+n_x} = y] \\ &= (P^{n_y})_{yx} (P^n)_{xx} (P^{n_x})_{xy} > 0. \end{aligned}$$

Therefore we have proven that for any $n \geq n_x + n_y + n_0$ one has $(P^n)_{yy} > 0$ which proves that y is aperiodic. The above reasoning holds for all $y \in \mathcal{X}$ so that P is aperiodic iff there exists $x \in \mathcal{X}$ aperiodic.

(ii) We proceed in similar fashion. Since P is irreducible there exists n_x such that $(P^{n_x})_{xy} > 0$ and since x is aperiodic there exists n_0 such that for all $n \geq n_0$ one has $(P^n)_{xx} > 0$. Then:

$$\begin{aligned} (P^{n+n_x})_{xy} &= \mathbb{P}_x[X_{n+n_x} = y] \\ &\geq \mathbb{P}_x[X_n = x, X_{n+n_x} = y] \\ &= (P^n)_{xx} (P^{n_x})_{xy} > 0, \end{aligned}$$

proving the second statement.

(iii) If P has a non-null diagonal there exists $x \in \mathcal{X}$ such that $P_{xx} > 0$, hence $(P^n)_{xx} \geq (P_{xx})^n > 0$ so that x is aperiodic and as a consequence P is aperiodic as well. \square

2.2.10 Convergence to the stationary distribution

We now show that aperiodicity is a sufficient condition for convergence to the stationary distribution. We recall the definition of the total variation distance, which we will use extensively:

Definition 2.2.11 Consider two probability distributions μ and μ' on \mathcal{X} . Then the total variation distance between μ and μ' is defined as:

$$\delta(\mu, \mu') = \max_{F \subset \mathcal{X}} \left| \sum_{x \in F} (\mu_x - \mu'_x) \right| = (1/2) \sum_x |\mu_x - \mu'_x|.$$

By definition we have $\delta(\mu, \mu') \in [0, 1]$.

The proof argument does not rely on the Perron-Frobenius theorem, which is noteworthy. The proof argument relies on a general technique called coupling. The reader may look at the discussion on mixing times for more information on coupling. The coupling used here can be summarized as follows. Choose an arbitrary state x , and imagine that one simulates two independent copies $(Y_n)_n$ and $(X_n)_n$ of the Markov chain. Assume that $(Y_n)_n$ starts with the stationary distribution and $(X_n)_n$ starts with an arbitrary distribution. Furthermore define T the first time that both chains are in state x . Then, define a third chain $Z_n = X_n \mathbf{1}\{n \leq T\} + Y_n \mathbf{1}\{n > T\}$ where we swap one for the other at time T . Then $(Z_n)_n$ has the same distribution as $(X_n)_n$, and since T is finite a.s. (as a consequence of aperiodicity), when $n \rightarrow \infty$ the limiting distribution of $(Z_n)_n$ must be that of $(Y_n)_n$, which is precisely the stationary distribution.

Theorem 2.2.12 Consider P an irreducible, positive recurrent and aperiodic transition matrix with unique stationary distribution π . Consider μ an arbitrary distribution and $(X_n)_n$ the (μ, P) -Markov chain.

(i) Then the distribution of X_n converges to π in total variation distance:

$$\sum_x |\mathbb{P}[X_n = x] - \pi_x| \xrightarrow{n \rightarrow \infty} 0.$$

(ii) The convergence rate is upper bounded by the inequality, for all n :

$$\sum_x |\mathbb{P}[X_n = x] - \pi_x| \leq \frac{C_{\mu, P}}{n}.$$

where $C_{\mu, P} > 0$ depends on initial distribution μ and transition matrix P , but is independent of n .

Proof. Define $(X_n, Y_n)_n$ a stochastic process on \mathcal{X}^2 such that:

- $(X_n)_n$ and $(Y_n)_n$ are independent,
- $(X_n)_n$ is the (μ, P) Markov chain,
- $(Y_n)_n$ is the (π, P) Markov chain.

Such a process indeed exists, as it suffices to choose $(X_n, Y_n)_n$ as a Markov chain with state space \mathcal{X}^2 , and transition matrix \tilde{P} expressed as follows:

$$\tilde{P}_{(x_1, x_2), (y_1, y_2)} = \mathbb{P}[(X_{n+1}, Y_{n+1}) = (x_2, y_2) | (X_n, Y_n) = (x_1, y_1)] = P_{x_1 x_2} P_{y_1 y_2}$$

Consider $x \in \mathcal{X}$ and define T the first time at which both chains are equal to x :

$$T = \min\{n \geq 0 : (X_n, Y_n) = x\}.$$

Irreducibility

Since P is aperiodic, from proposition 9 (second statement) for all $(x_1, x_2), (y_1, y_2)$ there exists $n \geq 0$ such that both $(P^n)_{x_1 x_2} > 0$ and $(P^n)_{y_1 y_2} > 0$. In turn:

$$(\tilde{P}^n)_{(x_1, x_2), (y_1, y_2)} = (P^n)_{x_1 x_2} (P^n)_{y_1 y_2} > 0,$$

therefore \tilde{P} is irreducible.

Positive recurrence

Since P is irreducible and positive recurrent, it has a unique stationary probability π . In turn, this implies that \tilde{P} has the unique stationary probability $(\pi_x \pi_y)_{x, y}$. Therefore \tilde{P} is positive recurrent. Since T is the first hitting time of state (x, x) for the Markov chain $(X_n, Y_n)_n$, we must have $T < \infty$ a.s.

Strong Markov property

T is a stopping time, T is finite a.s., and one has that $(X_T, Y_T) = (x, x)$ a.s. Therefore, by the strong Markov property, we must have that $(X_{T+n})_n$ and $(Y_{T+n})_n$ are two independent $(\delta(x), P)$ Markov chains, and are both independent of \mathcal{F}_T . Therefore $(Z_n)_n$ is a (μ, P) Markov chain. Hence, for all y :

$$\begin{aligned} \mathbb{P}[X_n = y] &= \mathbb{P}[Z_n = y] \\ &= \mathbb{P}[X_n = y, n \leq T] + \mathbb{P}[Y_n = y, n > T] \\ &= \mathbb{P}[X_n = y, n \leq T] + \mathbb{P}[Y_n = y] - \mathbb{P}[Y_n = y, n \leq T] \end{aligned}$$

Since $\mathbb{P}[Y_n = y] = \pi_y$ we have:

$$\left| \mathbb{P}[X_n = y] - \pi_y \right| \leq \mathbb{P}[X_n = y, n \leq T] + \mathbb{P}[Y_n = y, n \leq T].$$

Summing over y :

$$\sum_y \left| \mathbb{P}[X_n = y] - \pi_y \right| \leq 2\mathbb{P}[n \leq T] \xrightarrow{n \rightarrow \infty} 0,$$

since $T < \infty$ a.s., which concludes the proof of statement (i).

Applying the Markov inequality to the above we get:

$$\sum_y \left| \mathbb{P}[X_n = y] - \pi_y \right| \leq 2\mathbb{P}[n \leq T] \leq \frac{2\mathbb{E}[T]}{n}$$

which proves statement (ii) since we know that T has finite expectation. It is noted that, of course $\mathbb{E}[T]$ depends on the initial distribution μ , and the transition matrix P . □

2.3 Reversibility

In this section, we discuss a particular class of Markov chains, known as reversible Markov chains. Simply said, a Markov chain is reversible iff it is impossible to distinguish it (statistically) from the same Markov chain when time runs backwards. As we should see, not all Markov chains are reversible. Reversibility is a very interesting feature for many reasons, to name but a few:

- A reversible Markov chain is stationary
- When a Markov chain is reversible, it is possible to calculate its stationary distribution up to a normalization constant by inspection of the transition matrix. This is one of the reasons why many tractable Markovian models are reversible.
- Given an arbitrary distribution μ , it is possible to construct a reversible Markov chain whose stationary distribution is precisely μ . Furthermore, the constructed Markov chain is usually easy to sample from by simulation. This forms the basis for a large family of simulation algorithms known as MCMC (Markov Chain Monte Carlo).
- The mixing time of a Markov chain is easier to analyse when the Markov chain is reversible. The mixing time is the time required for a Markov chain to have a distribution close to the stationary distribution, when the initial distribution of the chain is arbitrary (see section 3.2 for a full discussion). Mixing time is one of the fundamental metrics to quantify the convergence speed of MCMC algorithms.

2.3.1 Reversible Markov chains

In this section we will consider Markov chains where the time index n is in \mathbb{Z} rather than in \mathbb{N} , mainly for ease of notation.

Definition 2.3.1 A Markov chain (X_n) is reversible iff for all $N \in \mathbb{Z}$, $(X_n)_n \stackrel{(d)}{=} (X_{N-n})_n$.

The first consequence of this definition is that a reversible Markov chain must be in stationary state.

Proposition 10 *A reversible Markov chain is stationary.*

Proof. Consider $(X_n)_n$ reversible. Consider $N \in \mathbb{Z}$. By definition $(X_n)_n \stackrel{(d)}{=} (X_{N-n})_n$, which implies that $(X_{N-n})_n$ is reversible as well. Therefore $(X_{N-n})_n \stackrel{(d)}{=} (X_{2N-(N-n)})_n = (X_{N+n})_n$. Hence for all $N \in \mathbb{Z}$, $(X_n)_n \stackrel{(d)}{=} (X_{n+N})_n$, so that $(X_n)_n$ is stationary. \square

2.3.2 Detailed balance conditions

If a Markov chain is reversible, it is stationary, so that its distribution must obey the full balance conditions. In fact reversibility implies that a much simpler set of equations, known as detailed balance conditions, must be satisfied.

Proposition 11 (*detailed balance conditions*) *Consider $(X_n)_n$ the (μ, P) -Markov chain. $(X_n)_n$ is reversible iff for all $x, y \in \mathcal{X}^2$:*

$$\mu_x P_{xy} = \mu_y P_{yx}.$$

Proof. Assume that $(X_n)_n$ is reversible. Then $(X_n)_n \stackrel{(d)}{=} (X_{1-n})_n$. As a consequence, for all $x, y \in \mathcal{X}^2$ we have $\mathbb{P}[X_0 = x, X_1 = y] = \mathbb{P}[X_0 = y, X_1 = x]$. By the Markov property: $\mathbb{P}[X_0 = x, X_1 = y] = \mu_x P_{xy}$, and by symmetry $\mathbb{P}[X_0 = x, X_1 = y] = \mu_x P_{xy}$ so that:

$$\mu_x P_{xy} = \mu_y P_{yx}.$$

On the other hand, assume that the detailed balance conditions hold. Consider $N \in \mathbb{Z}$, and $x_0, \dots, x_N \in \mathcal{X}^{N+1}$, by the Markov property, we have:

$$\begin{aligned} \mathbb{P}[X_0 = x_0, \dots, X_N = x_N] &= \mu_{x_0} P_{x_0, x_1} \dots P_{x_{N-1}, x_N} \\ \mathbb{P}[X_0 = x_N, \dots, X_N = x_0] &= \mu_{x_N} P_{x_1, x_0} \dots P_{x_N, x_{N-1}}. \end{aligned}$$

By detailed balance we have that, for all k :

$$\frac{P_{x_k, x_{k+1}}}{P_{x_{k+1}, x_k}} = \frac{\mu_{x_{k+1}}}{\mu_{x_k}},$$

so that dividing the two previous equations we get:

$$\begin{aligned} \frac{\mathbb{P}[X_0 = x_0, \dots, X_N = x_N]}{\mathbb{P}[X_0 = x_N, \dots, X_N = x_0]} &= \frac{\mu_{x_0}}{\mu_{x_N}} \prod_{k=0}^{N-1} \frac{P_{x_k, x_{k+1}}}{P_{x_{k+1}, x_k}} \\ &= \frac{\mu_{x_0}}{\mu_{x_N}} \prod_{k=0}^{N-1} \frac{\mu_{x_{k+1}}}{\mu_{x_k}} \\ &= 1. \end{aligned}$$

Therefore: $\mathbb{P}[X_0 = x_0, \dots, X_N = x_N] = \mathbb{P}[X_0 = x_N, \dots, X_N = x_0]$ for all N and all $(x_0, \dots, x_N) \in \mathcal{X}^{N+1}$, so that $(X_n)_n$ is indeed reversible. \square

In fact, consider the graph associated to the Markov chain. If this graph represents an electrical network, with edges representing power lines, there is a simple analogy between the detailed balance equations and the balance of electrical current over each line.

2.3.3 Examples and counter-examples

Let us first give a simple counter-example, so that not all Markov chains are not reversible, even if they are both aperiodic and irreducible. Consider the Markov chain on $\mathcal{X} = \{1, 2, 3\}$ with $0 < a < 1$:

$$P = \begin{pmatrix} (1-a) & 0 & a \\ a & (1-a) & 0 \\ 0 & a & (1-a) \end{pmatrix}$$

Exercise:

1. Draw the associated graph and give an intuitive explanation why such a Markov chain may/may not be reversible
2. Prove that P is irreducible and aperiodic
3. Write the full balance equations to find its stationary distribution μ .
4. Prove that the (μ, P) Markov chain is not reversible

Let us now give a generic example of reversible Markov chains, known as birth-and-death process. A birth-and-death process is a Markov chain on \mathbb{N} with $|X_{n+1} - X_n| \leq 1$ a.s. Those processes are called birth-and-death since they model the evolution of a population where, in each time step there may be at most one birth and one death. Indeed, at time n if there is a birth and no death $X_{n+1} = X_n + 1$, if there is a death and no birth $X_{n+1} = X_n - 1$ and otherwise $X_{n+1} = X_n$.

The transition matrix is given by:

$$P_{xy} = \begin{cases} \lambda(x) & \text{if } y = x + 1 \\ \mu(x) & \text{if } y = x - 1 \\ 1 - (\lambda(x) + \mu(x)) & \text{if } y = x \\ 0 & \text{otherwise .} \end{cases}$$

with $\mu(0) = 0$ and $\lambda(x) + \mu(x) \leq 1$ for all x .

Exercise:

1. Write the full balance equations and deduce whether or not the chain is reversible

2. Determine a stability condition using Foster's criterion
3. Deduce the stationary distribution

Birth-and-death processes (or their continuous time equivalents) are ubiquitous in queueing theory.

2.3.4 Stationary distribution

As suggested by the previous example, reversibility allows to calculate the stationary distribution by inspection of the transition matrix, without having to solve the full balance equations. Solving the full balance amounts to inverting a matrix whose size is the size of state space, and is infeasible for large state spaces.

We will use the graph notation to simplify the exposition. Consider $G(P)$ the graph associated to transition matrix P as defined in 2.1.5. Given a path $p = x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_N$ of length N in G we define the following function:

$$F(p) = \prod_{n=0}^{N-1} \frac{P_{x_n x_{n+1}}}{P_{x_{n+1} x_n}}.$$

This function plays a fundamental role and in fact can serve to calculate the stationary distribution as shown by the next result.

Proposition 12 Consider $(X_n)_n$ the (μ, P) -Markov chain, and assume that it is reversible and irreducible. Then:

- (i) For any path p from x to y we have: $F(p) = \frac{\mu_y}{\mu_x}$
- (ii) Define $F_{x,y} = F(p)$ where p is any path going from x to y . Then the stationary distribution is given by:

$$\mu_x = \frac{F_{\bar{x},x}}{\sum_{x \in \mathcal{X}} F_{\bar{x},x}}.$$

with $\bar{x} \in \mathcal{X}$ arbitrary.

Proof.

First statement

Consider x, y fixed and p an arbitrary path going from x to y : $x = x_0 \rightarrow \dots \rightarrow x_N = y$. Detailed balance holds so that for all $0 \leq n < N$:

$$\mu_{x_{n+1}} = \mu_{x_n} \frac{P_{x_n x_{n+1}}}{P_{x_{n+1} x_n}}$$

By induction we obtain that:

$$\mu_{x_N} = \mu_{x_0} \prod_{n=0}^{N-1} \frac{P_{x_n x_{n+1}}}{P_{x_{n+1} x_n}} = \mu_{x_0} F(p).$$

so that $F(p) = \frac{\mu_x N}{\mu_{x_0}} = \frac{\mu_y}{\mu_x}$.

Second statement

Consider $\bar{x} \in \mathcal{X}$ fixed. For all x we have that

$$\mu_x = \mu_{\bar{x}} F_{\bar{x}x}.$$

Furthermore $\sum_x \mu_x = 1$ so that by summing the above equation we get:

$$\mu_{\bar{x}} = \left(\sum_{x \in \mathcal{X}} F_{\bar{x}x} \right)^{-1},$$

so that for all x :

$$\mu_x = \frac{F_{\bar{x},x}}{\sum_{x \in \mathcal{X}} F_{\bar{x},x}}.$$

□

Indeed, the stationary distribution can be calculated by choosing an arbitrary state \bar{x} , and for each $x \in \mathcal{X}$ a path between \bar{x} and x . Since the choice of paths is arbitrary, we may choose it intelligently in order to minimize the computational effort.

2.3.5 A sufficient condition for reversibility

We give a condition for reversibility, based on inspection of the transition matrix. This criterion is known as Kolmogorov's criterion. We recall that a closed path is a path whose start and end vertices are identical.

Proposition 13 (*Kolmogorov's criterion*) *The (μ, P) Markov chain is reversible iff $F(p) = 1$ for any closed path p .*

Proof. Assume that the chain is reversible. We have seen that for any path p from x to y , we have $F(p) = \mu_y / \mu_x$. If p is closed, $x = y$ so that $F(p) = 1$.

Now assume that $F(p) = 1$ for any closed path. Consider x and y fixed, and p^1 and p^2 two paths from x to y . Define p^3 the path obtained by going from x to y using path p^1 and then going back from y to x using path p^2 (backwards). Since p^3 is a closed walk $F(p^3) = 1$. Furthermore, by definition of F : $F(p^3) = F(p^1) / F(p^2)$ so that $F(p^1) = F(p^2)$. Now choose \bar{x} arbitrary, and define $F_{\bar{x},x} = F(p)$ with p an arbitrary path from \bar{x} to x . By the previous reasoning, we may choose any path from x to y .

Then define the probability distribution: $\pi_x = F_{\bar{x}x} (\sum_x F_{\bar{x}x})^{-1}$. Consider y a neighbor of x and p^1 a path from \bar{x} to x . Construct a path p^2 from \bar{x} to y by first following p^1 and then going from x to y . We have that $F(p^2) = \frac{P_{xy}}{P_{yx}} F(p^1)$, so that: $F_{\bar{x}y} = \frac{P_{xy}}{P_{yx}} F_{\bar{x}x}$ and:

$$\pi_x P_{xy} = \pi_y P_{yx}.$$

Therefore π is the stationary distribution and it satisfies the detailed balance conditions so that reversibility holds, which concludes the proof. \square

This is a remarkable fact: we have seen that if a Markov chain is reversible, $F(p) = 1$ for any closed path p . Kolomogorov's criterion shows that this is also a necessary condition.

2.4 References

There are too many good treatments of Markov chains in the literature for us to cite them here. Sections 2.1 and 2.2 are based on [24]. Section 2.3 follows [17].

Chapter 3

Markov chains: stability and mixing Time

3.1 Stability and the Foster-Liapunov criterion

In practice, we are often presented with Markov chains which are not tractable, in the sense that it is not possible to calculate their stationary distribution (and related metrics) in closed form. For instance many controlled queuing systems used for modeling computer networks are not tractable. Therefore we settle for a much more modest goal: establish ergodicity of the Markov chain of interest. In particular, when the Markov chain describes a queuing system (the state of the Markov chain is the number of customers at the various servers), ergodicity corresponds to stability, namely the number of users waiting in the system (as a function of time) tends to a stationary process, and does not grow to infinity.

We have seen that a condition for ergodicity of a Markov chain is (i) irreducibility and (ii) positive recurrence of a given state. In practice establishing irreducibility is relatively easy, by inspection of the transition matrix. On the other hand, establishing positive recurrence of a state is not trivial, and in particular calculating $\mathbb{E}_x[T_x]$ explicitly is (most of the time) infeasible.

In this section we describe the Foster-Liapunov criterion which gives a simple condition for positive recurrence. It is the stochastic equivalent to the Liapunov condition for global stability of an equilibrium point of an ordinary differential equation (ODE). As its name indicates, it is due to F. G. Foster [9].

3.1.1 Rationale: the ODE case

The rationale for the Foster criterion is the well known Liapunov condition for (deterministic) ODEs. The Liapunov criterion [21] is ubiquitous in dynamical systems and

control theory, and is the main tool to analyse the asymptotic behaviour of solutions of ODEs.

Let us first recall the Liapunov criterion for ODEs. Proof is omitted, and can be found in any standard textbook on dynamical systems e.g. [29].

Theorem 3.1.1 (*Liapunov condition*) Consider the ODE:

$$\dot{x} = F(x).$$

with $F : \mathbb{R}^K \rightarrow \mathbb{R}^K$. Consider V a (Liapunov) function satisfying:

- (positive definiteness) $V(x) \geq 0$, with equality iff $x = 0$
- (negative drift) $\dot{V}(x) \leq 0$, with equality iff $x = 0$.
- (radial unboundedness) $V(x) \xrightarrow{x \rightarrow \infty} \infty$.

Then for all solutions x of the ODE we have $x(t) \xrightarrow{t \rightarrow \infty} 0$.

In fact, there is a natural analogy between Lipunov functions and potential energy of a physical system . Namely, potential energy is positive and tends to diminish as time increases, so that it is generally a good candidate when looking for Liapunov functions.

3.1.2 The Foster criterion

We return to Markov chains. First we define Liapunov functions in that setting. It should be noted that the definition chosen for Liapunov functions of Markov chains is a natural adaptation of the concept of Liapunov functions for ODE's.

Definition 3.1.2 Consider $V : \mathcal{X} \rightarrow \mathbb{R}$ and $(X_n)_n$ a Markov chain. V is a Liapunov function for the Markov chain $(X_n)_n$ iff there exists a finite set $\mathcal{Y} \subset \mathcal{X}$ and $\epsilon > 0$ such that:

- (positivity) $V(x) \geq 0$
- (finite expectation) $\mathbb{E}[V(X_{n+1})|X_n = x] < \infty$ for all $x \in \mathcal{Y}$.
- (negative drift) $\mathbb{E}[V(X_{n+1})|X_n = x] \leq V(x) - \epsilon$ for all $x \notin \mathcal{Y}$.

Remark 1 In matrix notation, identifying function V with the column vector $v = (V(x))_{x \in \mathcal{X}}$, the negative drift condition is simply written:

$$(Pv)_x \leq v_x - \epsilon, x \in \mathcal{Y}.$$

We may now state Foster's criterion for Markov chains.

Theorem 3.1.3 (*Foster's criterion for Markov chains*) Consider $(X_n)_n$ an irreducible Markov chain that admits a Liapunov function V . Define

$$T_{\mathcal{Y}} = \inf\{n \geq 0 : X_n \in \mathcal{Y}\},$$

the first hitting time of set \mathcal{Y} .

(i) Then, for all $x \in \mathcal{X}$ we have:

$$\mathbb{E}_x[T_{\mathcal{Y}}] \leq V(x)/\epsilon < \infty$$

(ii) Furthermore, $(X_n)_n$ is positive recurrent.

Proof. See subsection 3.1.5 □

3.1.3 Martingales

In order to facilitate the proof of the Foster criterion, we first introduce a result due to Doob called the optional stopping theorem. The optional stopping theorem allows to calculate the expected value of a certain type of processes (called martingales) evaluated at a stopping time. We first define martingales, sub-martingales and super-martingales. For a full discussion of martingales and optional stopping results, the interested reader may consult for instance [33].

Definition 3.1.4 Consider $(G_n)_n$ a stochastic process with values in \mathbb{R} , and define \mathcal{F}_n the σ -algebra generated by (G_0, \dots, G_n) .

- $(G_n)_n$ is a sub-martingale iff $\mathbb{E}[G_{n+1}|\mathcal{F}_n] \geq G_n$ for all n
- $(G_n)_n$ is a super-martingale iff $\mathbb{E}[G_{n+1}|\mathcal{F}_n] \leq G_n$ for all n
- $(G_n)_n$ is a martingale iff $\mathbb{E}[G_{n+1}|\mathcal{F}_n] = G_n$ for all n

Simply said, a sub-martingale (resp. super-martingale) is a process with negative (resp. positive) increments (in conditional expectation). A martingale is a process with increments that have null conditional expectation.

Let us give an example of a martingale. Consider a gambler playing repeatedly a game of head-or-tails, where the gambler must guess the outcome of a fair coin. If the gambler guesses correctly, he receives twice what he has bet, and otherwise the betted money is lost. Consider $(A_n)_n$ i.i.d Bernoulli with parameter $1/2$ representing the outcomes of trying to guess the outcomes of a fair coin at time instants $n \in \mathbb{N}$. At time n , the gambler has an amount of money G_n , bets an amount of money $B_n \geq 0$ which may depend both on (G_0, \dots, G_n) and (A_0, \dots, A_n) . He then gains an amount of money $(2A_{n+1} - 1)B_n$. The gambler's strategy is the choice of $(B_n)_n$, so that here we

allow for arbitrary adaptive strategies. Then regardless of the gambler's strategy, (G_n) is a martingale, since:

$$G_{n+1} = G_n + (2A_{n+1} - 1)B_n$$

so:

$$\mathbb{E}[G_{n+1}|\mathcal{F}_n] = G_n + B_n\mathbb{E}[(2A_{n+1} - 1)] = G_n$$

where we have used the fact that A_{n+1} is independent of \mathcal{F}_n , and that B_n is \mathcal{F}_n -measurable. As a consequence, for all n $\mathbb{E}[G_n] = G_0$ so that (in expectation), there exists no strategy that gains money after an arbitrary number of plays. Interestingly, there is no strategy that loses money either.

3.1.4 Optional stopping theorem

Theorem 3.1.5 (*Doob's optional stopping theorem*) Consider G_n a super-martingale with $\mathbb{E}[|G_0|] < \infty$ and T a stopping time. If we have either:

- (i) T is finite a.s. and $G_n \geq 0$ for all n ,
- (ii) T has finite expectation and there exists $g \geq 0$ such that for all n we have:
 $\mathbb{E}[|G_{n+1} - G_n||\mathcal{F}_n] \leq g$ a.s.

Then we have $\mathbb{E}[G_T] \leq \mathbb{E}[G_0]$.

Proof. Define the stopped process: $H_n = G_{n \wedge T}$. We have that:

$$H_{n+1} = H_n \mathbf{1}\{n \leq T\} + G_{n+1} \mathbf{1}\{n > T\}.$$

The events $\{n > T\}$ and $\{n \leq T\}$ are \mathcal{F}_n measurable since T is a stopping time. Taking conditional expectations:

$$\begin{aligned} \mathbb{E}[H_{n+1}|\mathcal{F}_n] &= H_n \mathbf{1}\{n \leq T\} + \mathbb{E}[G_{n+1}|\mathcal{F}_n] \mathbf{1}\{n > T\} \\ &\leq H_n \mathbf{1}\{n \leq T\} + G_n \mathbf{1}\{n > T\} \\ &= H_n (\mathbf{1}\{n \leq T\} + \mathbf{1}\{n > T\}) \\ &= H_n. \end{aligned}$$

Therefore $\mathbb{E}[H_{n+1}|\mathcal{F}_n] \leq H_n$ and $(H_n)_n$ is a super-martingale. Hence for all n :

$$\mathbb{E}[H_n] \leq \mathbb{E}[H_0].$$

Since $T < \infty$ a.s. we have that $H_n \xrightarrow[n \rightarrow \infty]{} G_T$ a.s. First assume that $\sup_n \mathbb{E}[|H_n|] < \infty$, then applying the dominated convergence theorem would yield the result:

$$\mathbb{E}[G_T] = \mathbb{E}[\lim_{n \rightarrow \infty} H_n] = \lim_{n \rightarrow \infty} \mathbb{E}[H_n] \leq \mathbb{E}[H_0].$$

Let us prove that we indeed have $\sup_n \mathbb{E}[|H_n|] < \infty$:

Case (i)

Since H_n is positive: $\mathbb{E}[|H_n|] = \mathbb{E}[H_n] \leq \mathbb{E}[H_0] = \mathbb{E}[G_0] < \infty$, so that $\sup_n \mathbb{E}[|H_n|] < \infty$.

Case (ii)

We have that:

$$|Y_{n+1}| \leq |Y_n| + |Y_{n+1} - Y_n|.$$

Furthermore $Y_{n+1} - Y_n = \mathbf{1}\{n > T\}(G_{n+1} - G_n)$. So taking conditional expectations:

$$\mathbb{E}[|Y_{n+1}| | \mathcal{F}_n] \leq |Y_n| + \mathbf{1}\{n > T\} \mathbb{E}[|G_{n+1} - G_n| | \mathcal{F}_n] \leq |Y_n| + g \mathbf{1}\{n > T\}.$$

Taking expectations:

$$\mathbb{E}[|Y_{n+1}|] \leq \mathbb{E}[|Y_n|] + g \mathbb{P}[n > T].$$

Summing the above inequality yields:

$$\sup_n \mathbb{E}[|Y_n|] \leq \mathbb{E}[|Y_0|] + g \sum_{n \geq 0} \mathbb{P}[n > T] = \mathbb{E}[|G_0|] + g \mathbb{E}[T] < \infty.$$

so $\sup_n \mathbb{E}[|Y_n|] < \infty$ which concludes the proof. \square

Going back to the gambler's example, Doob's optional stopping theorem shows that, even if the gambler is allowed to leave the game after a given random time (that might depend on his current fortune), there is still no way to earn money in expectation (poor gambler ...).

3.1.5 Proof of Foster's criterion

We are now fully equipped to prove Foster's criterion. It is interesting to see that the optional stopping theorem allows for a relatively short and elementary proof.

First statement

Let us first prove that $T_y < \infty$ a.s. Define $G_n = V(X_{n \wedge T_y}) + \epsilon(n \wedge T_y)$.

- If $n \geq T_y$, we have that $G_{n+1} = G_n$.
- If $n < T_y$ we have that $X_n \notin \mathcal{Y}$ and that $G_n = V(X_n) + n\epsilon$ and $G_{n+1} = V(X_{n+1}) + \epsilon(n+1)$.

Hence:

$$G_{n+1} - G_n = \mathbf{1}\{n < T_y\}(V(X_{n+1}) - V(X_n) + \epsilon).$$

Since event $\{n < T_y\} = \cap_{n' \leq n} \{n' \neq T_y\}$ is \mathcal{F}_n -measurable, we may take conditional expectations:

$$\mathbb{E}_x[G_{n+1} | \mathcal{F}_n] = G_n + \mathbf{1}\{n < T_y\}(\mathbb{E}_x[V(X_{n+1}) | \mathcal{F}_n] - V(X_n) + \epsilon) \leq G_n,$$

by the negative drift assumption and the fact that $n > T_{\mathcal{Y}}$ implies $X_n \notin \mathcal{Y}$. So $(G_n)_n$ is a super-martingale and for all n we have:

$$\mathbb{E}_x[G_n] \leq \mathbb{E}_x[G_0] = V(x)$$

Since V is positive we have that $\mathbf{1}\{n \leq T_{\mathcal{Y}}\}n\epsilon \leq G_n$. Taking expectations: $n\epsilon\mathbb{P}_x[n \leq T_{\mathcal{Y}}] \leq \mathbb{E}_x[G_n] \leq V(x)$. So for all n we have:

$$\mathbb{P}_x[n \leq T_{\mathcal{Y}}] \leq V(x)/(n\epsilon).$$

Letting $n \rightarrow \infty$ in the above expression, we have that $\mathbb{P}_x[n \leq T_{\mathcal{Y}}] \xrightarrow{n \rightarrow \infty} 0$ so that $T_{\mathcal{Y}} < \infty$ a.s.

Now, since $(G_n)_n$ is a positive super-martingale, and that $T_{\mathcal{Y}}$ is finite a.s., we may apply Doob's optional stopping theorem (first set of conditions) to yield $\mathbb{E}_x[G_{T_{\mathcal{Y}}}] \leq \mathbb{E}_x[G_0] = V(x)$ so that:

$$\epsilon\mathbb{E}_x[T_{\mathcal{Y}}] \leq \mathbb{E}_x[V(X_{T_{\mathcal{Y}}})] + \epsilon\mathbb{E}_x[T_{\mathcal{Y}}] = \mathbb{E}_x[G_{T_{\mathcal{Y}}}] \leq V(x)$$

and hence $\mathbb{E}_x[T_{\mathcal{Y}}] \leq V(x)/\epsilon < \infty$ which concludes the proof of the first statement.

Second statement

We will use the result of the first statement. We consider $x \in \mathcal{Y}$, and, in order to complete the proof it is sufficient to prove that $\mathbb{E}_x[T_x] < \infty$ with $T_x = \min\{n \geq 1 : X_n = x\}$ the return time to x .

To prove the result we are going to consider the values of the Markov chain $(X_n)_n$ at time instants n such that $X_n \in \mathcal{Y}$. It is noted that the times at which $X_n \in \mathcal{Y}$ are random.

We define the successive return times to \mathcal{Y} as follows: $T^{-1} = 0$, and for all $k \geq 0$:

$$T^k = \min\{n > T^{k-1} : X_n \in \mathcal{Y}\}.$$

Define $V_{\mathcal{Y}} = \max_{x \in \mathcal{Y}} V(x)$. By the first statement of the theorem $\mathbb{E}_x[T^{n+1} - T^n] \leq V_{\mathcal{Y}}/\epsilon$ for all n , so that in particular $T^k < \infty$ a.s. We define the Markov chain sampled at the return times to \mathcal{Y} : $Y_n = X_{T^n}$. We have that Y_n is a Markov chain on \mathcal{Y} , which is irreducible (left as an exercise). Define $\tilde{T}_x = \min\{n \geq 1 : Y_n = x\}$ the return time to x for this sampled chain. It is noted that, by definition, we have that $T_x = T^{\tilde{T}_x}$. Since \mathcal{Y} is finite, Y_n must be positive recurrent, so that $\mathbb{E}[\tilde{T}_x] < \infty$.

Define the process $G_n = T^n - nV_{\mathcal{Y}}/\epsilon$. From the strong Markov property, $T^{n+1} - T^n$ is independent of \mathcal{F}_{T^n} , so:

$$\begin{aligned} \mathbb{E}_x[G_{n+1} - G_n | G_n, \dots, G_0] &= \mathbb{E}[T^{n+1} - T^n | \mathcal{F}_{T^n}] - V_{\mathcal{Y}}/\epsilon \\ &= \mathbb{E}[T^{n+1} - T^n] - V_{\mathcal{Y}}/\epsilon \leq 0 \end{aligned}$$

so that G_n is a supermartingale. Furthermore:

$$\begin{aligned}\mathbb{E}_x[|G_{n+1} - G_n| | G_n, \dots, G_0] &\leq \mathbb{E}_x[T^{n+1} - T^n | \mathcal{F}_{T^n}] + V_Y/\epsilon \\ &= \mathbb{E}[T^{n+1} - T^n] + V_Y/\epsilon \\ &\leq 2V_Y/\epsilon < \infty,\end{aligned}$$

applying once again the first statement of the theorem. Therefore we may apply the optional stopping theorem (second set of conditions) to yield $\mathbb{E}_x[G_{\tilde{T}_x}] \leq \mathbb{E}_x[G_0] = 0$. Since $G_{\tilde{T}_x} = T^{\tilde{T}_x} - \tilde{T}_x V_Y/\epsilon = T_x - \tilde{T}_x V_Y/\epsilon$ we have proven:

$$\mathbb{E}_x[T_x] \leq \mathbb{E}_x[\tilde{T}_x] V_Y/\epsilon < \infty.$$

which concludes the proof. □

3.1.6 An illustration in one dimension

We propose to give a simple illustration of the Foster criterion for Markov chains on \mathbb{N} , i.e. a single server queue. Consider the following Markov chain on \mathbb{N} : $X_{n+1} = 0 \wedge (X_n + A_n - D_n)$, with A_n and D_n positive random variables with finite expectation. We assume that the distribution of A_n and D_n are independent of X_0, \dots, X_{n-1} , and n but might depend on X_n . This Markov chain represents for instance the number of customers in a queue, with A_n the number of customers arriving in time interval $[n, n + 1]$ and D_n the number of customers departing during time interval $[n, n + 1]$. Their distributions might depend on the current number of customers X_n , so that the speed at which users enter and leave the system is state-dependent.

Exercise:

1. Give a sufficient condition for irreducibility of X_n
2. Prove that $V(x) = x \wedge 0$ is a Liapunov function
3. Define $d(x) = \mathbb{E}[A_n - D_n | X_n = x]$, what is intuitive meaning of d ?
4. Give a stability condition for (X_n) as a function of d .

3.2 Mixing time of Markov chains

3.2.1 Sampling from a stationary distribution

Throughout this section we will be working with Markov chains on *finite* (but possibly large) state spaces.

Consider P a transition matrix which is irreducible, aperiodic and positive recurrent, with unique stationary distribution π . Assume that we would like to sample from π . We have previously proven (in Theorem 2.2.12) that for irreducible, aperiodic and positive

recurrent Markov chains, the distribution of $(X_n)_n$ converges to the stationary distribution π . Therefore the simplest approach would be to draw a sample path X_0, \dots, X_n , with an arbitrary starting distribution μ . If n is large enough, the distribution of X_n should be a good approximation to π . Drawing several independent sample paths would yield i.i.d copies of X_n , so that one may calculate any quantity of interest related to the stationary distribution π .

Now, it is clear that the main problem is the choice of n . If n is too small the distribution of X_n will be a poor approximation of π , and if n is too large the simulation procedure will require waiting for more time than necessary. Typically, given a tolerance threshold ϵ , and an initial distribution μ one would like to find $n_{\epsilon, \mu}$ such that for all $n \geq n_{\epsilon, \mu}$ the total variation distance between π and the distribution of X_n is at most ϵ . Theorem 2.2.12 (second statement) in fact gives us a first lead, so that choosing $n_{\epsilon, \mu} = O(1/\epsilon)$ would be fine. There are two problems here:

- The convergence rate (see the example below) prescribed by Theorem 2.2.12 is too slow, and under mild assumptions, the convergence speed to the stationary distribution is exponential. Hence we may choose $n_{\epsilon, \mu} = O(\log(1/\epsilon))$
- Theorem 2.2.12 does not specify the value of $C_{\mu, P}$, and $C_{\mu, P}$ is in general hard to calculate. Therefore we only know the proper value of $n_{\epsilon, \mu}$ up to a constant that depends on μ and P . We will show that it is possible to choose $n_{\epsilon, \mu}$ independently of μ (so that the convergence is uniform in terms of the initial distribution).

The threshold n_ϵ is sometimes called the "burn-in time" and represents the amount of time required by the method to obtain a sample of a distribution approximating π with an error of at most ϵ .

In fact being able to sample from the stationary distribution of certain transition matrices is at the heart of many important problems in applied mathematics, to cite but a few: global optimization of a function, image reconstruction, resource allocation in networks, sampling from posterior distributions... Algorithms that sample from the stationary distribution of a transition matrix in order to solve a problem are grouped under the generic name MCMC (Markov Chain Monte Carlo) and will be addressed in latter chapters.

As a side note, there are more sophisticated algorithms such as Propp and Wilson's algorithm, that are able to sample *exactly* from π . Such algorithms are named perfect simulation algorithms. Most of these algorithms draw a sample path *of random length* X_0, \dots, X_T where T is a well chosen stopping time. T is chosen such that X_T has exactly distribution π , and so that $\mathbb{E}[T]$ is as small as possible. The counterpart is that, in many models of interest T is so large that one has no choice but using the crude method described above.

To summarize, we are interested in upper bounding the *mixing time*, which is the convergence time to the stationary distribution. Furthermore, in practice we would like

to be able to have a tractable way to calculate the mixing time as a function of the transition matrix P .

3.2.2 Example for two states

In order to understand the problem in more details, we start by studying the simplistic case where there are 2 states. In this case we may in fact calculate the mixing time in closed form, illustrating the link between the mixing time and the eigenvalues of P .

We consider $\mathcal{X} = \{0, 1\}$, and the transition matrix:

$$P = \begin{pmatrix} 1-a & a \\ b & 1-b \end{pmatrix}.$$

with $(a, b) \in (0, 1)^2$. We may readily verify that P is irreducible, aperiodic and positive recurrent (by finiteness of the state space) with stationary distribution $\pi = (\frac{b}{a+b}, \frac{a}{a+b})$. Define $\mu(n)$ the distribution of X_n . Define the total variation distance between $\mu(n)$ and π :

$$\delta(\mu(n), \pi) = \frac{1}{2}(|\mu_0(n) - \pi_0| + |\mu_1(n) - \pi_1|) = |\mu_0(n) - \pi_0|.$$

It remains to calculate $\mu_0(n) - \pi_0$. We proceed by induction. The full balance equations give:

$$\begin{aligned} \mu_0(n+1) &= (1-a)\mu_0(n) + b\mu_1(n) \\ &= (1-a)\mu_0(n) + b(1-\mu_0(n)) \\ &= (1-a-b)\mu_0(n) + b. \end{aligned}$$

Furthermore:

$$\pi_0 = (1-a-b)\pi_0 + b.$$

Subtracting the two equations above we get:

$$\mu_0(n+1) - \pi_0 = (1-a-b)(\mu_0(n) - \pi_0)$$

Hence $\delta(\mu(n+1), \pi) = |1-a-b|\delta(\mu(n), \pi)$ and by induction we have calculated the convergence time exactly:

$$\delta(\mu(n), \pi) = |1-a-b|^n \delta(\mu(0), \pi).$$

Several remarks can be made:

- The convergence to the stationary distribution is exponential with rate $|1-a-b|$ (the mixing time)

- The total variation distance is upper bounded by 1, so that $\delta(\mu(n), \pi) \leq |1 - a - b|^n$, irrespective of the starting distribution $\mu(0)$. This shows that convergence to the stationary distribution is *uniform* in the initial distribution.
- The fastest mixing is obtained when $a + b = 1$. Here the chain is in stationary state after *one transition* so that $\mu(n) = \pi, n \geq 1$.
- Mixing can be arbitrary slow when $a + b$ is close to 0. In this case the chain tends to stay in the same state for long periods of times. A clearer link between mixing and the probability of escaping sets of states will be given later.
- The eigenvalues of P are 1 and $1 - a - b$, so that the mixing time is exactly the absolute value of the second largest eigenvalue of P . This is true for larger states spaces as well, see below.

3.2.3 Exponential Mixing

We start by proving that any irreducible, positive recurrent and aperiodic Markov chain has exponential mixing. Namely the total variation distance between the distribution of X_n and the stationary distribution as a function of n decreases exponentially. The proof is constructive so that we calculate an upper bound for the mixing time explicitly. The proof is short by relying on an interesting trick.

Theorem 3.2.1 *Consider \mathcal{X} finite. Consider P an irreducible, positive recurrent and aperiodic transition matrix with unique stationary probability π . Consider $(X_n)_n$ the (μ, P) Markov chain, and define $\mu(n)$ the distribution of X_n . Then:*

- (i) *There exists $C_P > 0$ and $\theta_P > 0$ such that for all n and all $\mu(0)$:*

$$\delta(\mu(n), \pi) \leq C_P \theta_P^n$$

- (ii) *Consider r such that $\min_{xy} (P^r)_{xy} > 0$. Then we may choose:*

$$\begin{aligned} C_P &= (1 - d)^{-1}, \\ \theta &= (1 - d)^{1/r}, \\ d &= \min_{x,y} \frac{(P^r)_{xy}}{\pi_y}. \end{aligned}$$

Proof.

If P has only non-null entries:

First assume that $\min_{xy} (P)_{xy} > 0$. Define $d = \min_{x,y} P_{xy}/\pi_y$. Define the transition matrix \tilde{P} :

$$\tilde{P}_{xy} = \frac{P_{xy} - d\pi_y}{1 - d}$$

One may readily check that \tilde{P} is a transition matrix. Consider $(A_n)_n$ i.i.d. Bernoulli with parameter d . Define $(Y_n)_n$ the Markov chain such that Y_0 has distribution μ , and:

$$\mathbb{P}[Y_{n+1} = y | Y_n = x, A_n] = \mathbf{1}\{A_n = 1\}\pi_y + \mathbf{1}\{A_n = 0\}\tilde{P}_{xy}.$$

Namely, $(Y_n)_n$ a Markov chain that may be simulated the following way: at time n , one first draws A_n a Bernoulli variable with parameter d . Then, if $A_n = 0$, Y_{n+1} is drawn with distribution $(\tilde{P}_{Y_n y})_y$, and if $A_n = 1$, Y_{n+1} is drawn with distribution π , irrespective of the value of Y_n . We have that (Y_n) is the (μ, P) Markov chain since:

$$\begin{aligned} \mathbb{P}[Y_{n+1} = y | Y_n = x] &= \mathbb{P}[A_n = 1]\pi_y + \mathbb{P}[A_n = 0]\tilde{P}_{xy} \\ &= d\pi_y + (1 - d)\frac{P_{xy} - d\pi_y}{1 - d} \\ &= P_{xy}. \end{aligned}$$

The distribution of Y_n conditional to A_1, \dots, A_n is given by:

$$Y_n \stackrel{(d)}{=} \begin{cases} \pi, & \text{if } \sum_{n'=1}^n A_{n'} > 0 \\ \mu\tilde{P}^n & \text{if } \sum_{n'=1}^n A_{n'} = 0. \end{cases}$$

Using the two facts:

- Y_n has the same distribution as X_n
- the total variation distance between two arbitrary distributions is bounded by 1

we have that:

$$\begin{aligned} \delta(\mu(n), \pi) &= \mathbb{P}\left[\sum_{n'=1}^n A_{n'} > 0\right] \delta(\pi, \pi) + \mathbb{P}\left[\sum_{n'=1}^n A_{n'} = 0\right] \delta(\mu\tilde{P}^n, \pi) \\ &\leq \mathbb{P}\left[\sum_{n'=1}^n A_{n'} = 0\right] = (1 - d)^n \end{aligned}$$

General case:

Since P is aperiodic and \mathcal{X} is finite, there exists r such that $\min_{xy}(P^r)_{xy} > 0$. Consider $j \in \{0, \dots, r - 1\}$ and define $(Z_n)_n = (X_{kn+j})_n$. $(Z_n)_n$ is a Markov chain with initial distribution μP^j and transition matrix P^r . Since $\min_{xy}(P^r)_{xy} > 0$ we can define $d = \min_{x,y} \frac{(P^r)_{xy}}{\pi_y}$ and apply the reasoning above to yield, for all $m \geq 0$:

$$\delta(\mu(rm + j), \pi) \leq (1 - d)^m.$$

Consider $n \geq 0$ and write n as: $n = rm + j$, so that $m = (n - j)/r > (n - r)/r = n/r - 1$, then the above inequality reads:

$$\delta(\mu(n), \pi) \leq (1 - d)^{n/r-1} = C_P \theta_P^n,$$

by defining $C_P = (1 - d)^{-1}$ and $\theta_P = (1 - d)^{1/r}$. The above inequality holds for all n , so that the proof is complete. □

3.2.4 Mixing Time

We may now define the mixing time of an ergodic, aperiodic Markov chain. Mixing time is simply defined as the minimal n such that, when the distribution of X_0 is arbitrary, the total variation distance between the stationary distribution and the distribution of X_n is less than $1/4$.

Definition 3.2.2 Consider $(X_n)_n$ the (μ, P) Markov chain. Assume that P is aperiodic, irreducible and positive recurrent with unique stationary distribution π . Define:

$$d(n) = \max_x \delta((P^n)_{x,\cdot}, \pi).$$

The mixing time τ is defined as:

$$\tau = \min\{n \geq 0 : d(n) \leq 1/4\}.$$

An observation on the definition of mixing time can be made: $(P^n)_{x,\cdot}$ is the distribution of X_n conditional to $X_0 = x$. Therefore, $d(n)$ is supremum of the total variation distance between the distribution of X_n and the stationary distribution π , where the supremum is taken over all possible starting states.

3.2.5 Standardizing distance

The value of $1/4$ might at first seem arbitrary, but it is not the case, as shown below. Essentially, every τ iterations, the distance between the distribution of X_n and the stationary distribution is halved. We also introduce a function related to d , written:

$$\bar{d}(n) = \sup_{x,y} \delta((P^n)_{x,\cdot}, (P^n)_{y,\cdot}),$$

which quantifies the maximal total variation distance between the distributions of $(X_n)_n$ for different starting states.

Proposition 14 *For all n one has:*

$$d(n) \leq \bar{d}(n) \leq 2d(n).$$

Furthermore, \bar{d} is sub-multiplicative, ie for all n, n' :

$$\bar{d}(n + n') \leq \bar{d}(n)\bar{d}(n').$$

Proof. ... □

Corollary 3.2.3 *Consider $n \geq 0$, then one has that $d(n\tau) \leq 2^{-n}$.*

Proof. We have that:

$$d(n\tau) \leq \bar{d}(n\tau) \leq (\bar{d}(\tau))^n.$$

by proposition 14. Furthermore:

$$\bar{d}(\tau) \leq 2d(\tau) \leq 1/2.$$

hence: $d(n\tau) \leq 2^{-n}$ as announced. □

3.2.6 Coupling

Here we introduce a general purpose technique for upper bounding the mixing time of Markov chains called coupling. The basic idea is to consider two sample paths of the same Markov chain starting in two different states x and y , and to define the first time $\tau_{x,y}$ that those chains meet at some state. Then the distribution of $\tau_{x,y}$ allows to upper bound the mixing time.

In fact, the reader might notice that, in order to prove Theorem 2.2.12 (convergence to the stationary distribution), we relied on the same argument, namely we considered two chains, one starting with the stationary distribution, and another one starting with an arbitrary distribution.

Proposition 15 *Assume that for all $x, y \in \mathcal{X}$ there exists a stochastic process $(X_n^{x,y}, Y_n^{x,y})_n$ such that $X_0^{x,y} = x, Y_0^{x,y} = y$ and $(X_n^{x,y})_n, (Y_n^{x,y})_n$ are both Markov chains with transition matrix P . Define $\tau_{x,y} = \min\{n \geq 0 : X_n^{x,y} = Y_n^{x,y}\}$.*

(i) *Then one has:*

$$d(n) \leq \max_{x,y} \mathbb{P}[\tau_{x,y} > n].$$

(ii) *By corollary, the mixing time is upper bounded as:*

$$\tau \leq 4 \max_{x,y} \mathbb{E}[\tau_{x,y}].$$

3.3 References

Section 3.2 is based on [20], which covers the topic of mixing time in details. Section 3.1 is based on [7]. A through treatment of stability for Markov processes and Foster-Liapunov-like results is found in [23].

Chapter 4

Stochastic approximation

4.1 The basic stochastic approximation scheme

4.1.1 A first example

We propose to start the exposition of the topic by an example. The arguments are given in a crude manner. Formal proofs will be given in section 4.2. This example is taken from the very article [26] which introduced stochastic approximation. Consider $x \in \mathbb{R}$ the parameter of a system and $g(x) \in \mathbb{R}$ an output value from this system when parameter x is used. We assume g to be a smooth, increasing function. An agent wants to determine sequentially $x^* \in \mathbb{R}$ the value such that the system output equals a target value g^* . If for all x , the value of $g(x)$ can be observed directly from the system, then determining g^* could be solved by a simple search technique such as binary search or golden ratio search. Here we assume that only a *noisy* version of g can be observed. Namely, at time $n \in \mathbb{N}$, the decision maker sets the parameter equal to x_n , and observes $Y_n = g(x_n) + M_n$ with M_n a random variable denoting noise, with $\mathbb{E}[M_n] = 0$.

In order to determine $g(x)$, a crude approach would be to sample parameter x repeatedly and average the result, so that the effect of noise would cancel out, and apply a deterministic line search (such as binary search).

[26] proposed a much more elegant approach. If $x_n > x^*$, we have that $g(x_n) > g^*$, so that diminishing x_n by a small amount proportional to $g^* - g(x_n)$ would guarantee $x_{n+1} \in [x^*, x_n]$. Therefore, define ϵ_n a sequence of small positive numbers, and consider the following update scheme:

$$x_{n+1} = x_n + \epsilon_n(g^* - Y_n) = x_n + \epsilon_n(g^* - g(x_n)) + \epsilon_n M_n.$$

The first intuition is that if the noise sequence is well behaved (say $\{M_n\}$ is i.i.d Gaussian with mean 0 and variance 1) and $\epsilon_n = 1/n$, then the law of large numbers would guarantee that the noise “averages out” so that for large n noise can be

ignored altogether. Namely, define $S_n = \sum_{k \geq n} M_k/k$, then $\text{var}(S_n)$ is upper bounded by $\sum_{k \geq n} 1/k^2 \rightarrow_{n \rightarrow +\infty} 0$, so that S_n should be negligible. (Obviously this reasoning is heuristic and to make it precise we have to use a law of large numbers-like result ...)

Now assume no noise ($M_n \equiv 0$ for all n), and $\epsilon_n = 1/n$, g smooth with a strictly positive first derivative upper bounded by g' . Removing the noise term M_n :

$$g(x_{n+1}) = g(x_n) + \frac{g'(x_n)}{n}(g^* - g(x_n)).$$

By the fundamental theorem of calculus: $(1/n)|g^* - g(x_n)| \leq (g'/n)|x^* - x_n|$. So for $n \geq g'$, we have either $x_n \leq x_{n+1} \leq x^*$ or $x_n \geq x_{n+1} \geq x^*$. In both cases, $n \mapsto |g(x_n) - g^*|$ is decreasing for large n .

It is also noted that:

$$\frac{x_{n+1} - x_n}{n} = (g^* - g(x_n)),$$

so that x_n appears as a discretization (with discretization steps $\{1/n\}$) of the following o.d.e.:

$$\dot{x} = g^* - g(x).$$

This analogy will be made precise in the next subsection.

4.1.2 The associated o.d.e

We now introduce the so-called o.d.e. approach popularized by [22], which allow to analyze stochastic recursive algorithms such as the one considered by Robbins in his original paper. It is noted that [26] did not rely on the o.d.e. method and used direct probabilistic arguments.

The crude reasoning above suggests that the asymptotic behavior of the random sequence $\{x_n\}$ can be obtained by determining the asymptotic behavior of a corresponding (deterministic) o.d.e. In this lecture we will consider a sequence $x_n \in \mathbb{R}^d$, $d \geq 1$, and a general update scheme of the following form:

$$x_{n+1} = x_n + \epsilon_n(h(x_n) + M_n),$$

with $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$. We define the associated o.d.e. :

$$\dot{x} = h(x).$$

We will prove that (with suitable assumptions on h , the noise and step sizes) if the o.d.e. admits a continuously differentiable Liapunov function V , then we have that $V(x_n) \rightarrow_{n \rightarrow +\infty} 0$ almost surely. We recall that V is a Liapunov function if it is positive, radially unbounded, and strictly diminishing along the solutions of the o.d.e.

4.1.3 Instances of stochastic approximation algorithms

Algorithms based on stochastic approximation schemes have become ubiquitous in various fields, including signal processing, optimization, machine learning and economics/game theory. There are several reasons for this:

- **Low memory requirements:** the basic stochastic approximation is a Markovian update: the value of x_{n+1} is a function of x_n and the observation at time n . So its implementation requires a small amount of memory.
- **Influence of noise:** stochastic approximation algorithms are able to work with noise, so that they are good candidates as “on-line” optimization algorithms which work with the noisy output of a running system. Furthermore the convergence of a stochastic approximation scheme is determined by inspecting a deterministic o.d.e. which is simpler to analyze and does not depend on the statistics of the noise.
- **Iterative updates:** Once again since they are Markovian updates, stochastic approximation schemes are good models for collective learning phenomena where a set of agents interact repeatedly and update their behavior depending on their most recent observation. This is the reason why results on learning schemes in game theory rely heavily on stochastic approximation arguments.

We give a few examples of stochastic algorithms found in the literature.

4.1.4 Stochastic gradient algorithms

Stochastic gradient algorithms allow to find a local maximum of a cost function whose value is only known through noisy measurements, and are commonplace in machine learning (on-line regression, training of neural networks, on-line optimization of Markov decision processes etc). We consider a function $f : \mathbb{R} \rightarrow \mathbb{R}$ which is strongly convex, twice differentiable with a unique minimum x^* . f cannot be observed directly, nor can its gradient ∇f . At time n we can observe $f(x_n) + M_n$. Therefore it makes sense to approximate ∇f by finite differences, with a suitable discretization step. Consider the scheme (due to Kiefer and Wolfowitz [18]):

$$x_{n+1} = x_n - \epsilon_n \frac{f(x_n + \delta_n) - f(x_n - \delta_n)}{2\delta_n},$$

The associated o.d.e is $\dot{x} = -\nabla f(x)$ which admits the Liapunov function $V(x) = f(x) - f(x^*)$. With the proper step sizes (say $\epsilon_n = n^{-1}$, $\delta_n = n^{-1/3}$) it can be proven that the method converges to the minimum $x_n \xrightarrow{n \rightarrow \infty} f(x^*)$ almost surely.

4.1.5 Distributed updates

In many applications, the components of x_n are not updated simultaneously. This is for instance the case in distributed optimization when each component of x_n is controlled by a different agent. This is also the case for on-line learning algorithms for Markov Decision Processes such as Q-learning.

For instance assume that at time n , a component $k(n)$ uniformly distributed $\{1, \dots, d\}$ is chosen, and only the $k(n)$ -th component of x_n is updated:

$$x_{n+1,k} = \begin{cases} x_{n,k} + \epsilon_n(h_k(x_n) + M_{n,k}) & , k = k(n) \\ x_{n,k} & , k \neq k(n) \end{cases}.$$

Then it can be proven that the behavior of $\{x_n\}$ can be described by the o.d.e. $\dot{x} = h(x)$. Namely, the asymptotic behavior of $\{x_n\}$ is the same as in the case where all its components are updated simultaneously. This is described for instance in [5][Chap 7].

4.1.6 Fictitious play

Fictitious play is a learning dynamic for games introduced by [8], and studied extensively by game theorists afterwards (see for instance [10]). Consider 2 agents playing a matrix game. Namely, at time $n \in \mathbb{N}$, agent $k \in \{1, 2\}$ chooses action $a_n^k \in \{1, \dots, \mathcal{A}\}$, and receives a reward A_{a_1, a_2}^k , where A^1, A^2 are two \mathcal{A} by \mathcal{A} matrices with real entries. Define the empirical distribution of actions of player k at time n by :

$$p^k(a, n) = \frac{1}{n} \sum_{t=1}^n \mathbf{1}\{a_t^k = a\}.$$

A natural learning scheme for agent k is to assume that at time $n + 1$, agent k' will choose an action whose probability distribution is equal to $p^{k'}(\cdot, n)$, and play the best action. Namely agent k assumes that $\mathbb{P}[a_{n+1}^{k'} = a] = p^{k'}(a, n)$, and chooses the action maximizing his expected payoff given that assumption.

We define

$$g^k(\cdot, p') = \max_{p \in \mathcal{P}} \sum_{1 \leq a \leq \mathcal{A}} \sum_{1 \leq a' \leq \mathcal{A}} p(a) A_{a, a'}^k p'(a'),$$

with \mathcal{P} the set of probability distributions on $\{1, \dots, \mathcal{A}\}$. $g^k(\cdot, p')$ is the probability distribution of the action of k maximizing the expected payoff, knowing that player k' will play an action distributed as p' .

The empirical probabilities can be written recursively as:

$$(n + 1)p^k(a, n + 1) = np^k(a, n) + \mathbf{1}\{a_n^k = a\},$$

so that:

$$p^k(a, n+1) = p^k(a, n) + \frac{1}{n+1}(\mathbf{1}\{a_n^k = a\} - p^k(a, n)).$$

Using the fact that $\mathbb{E}[\mathbf{1}\{a_n^k = a\}] = g^k(\cdot, p^k)$, we recognize that the probabilities p are updated according to a stochastic approximation scheme with $\epsilon_n = 1/(n+1)$, and the corresponding o.d.e. is

$$\dot{p} = g(p) - p.$$

It is noted that such an o.d.e may have complicated dynamics and might not admit a Liapunov function without further assumptions on the structure of the game (the matrices A^1 and A^2).

4.2 Convergence to the o.d.e limit

In this section we prove the basic stochastic approximation convergence result for diminishing step sizes with martingale difference noise. This setup is sufficiently simple to grasp the proof techniques without relying on sophisticated results. The only prerequisites are: the (discrete-time) martingale convergence theorem and two basic results on o.d.e.'s, namely Gronwall's inequality and the Picard - Lindelöf theorem. We largely follow the exposition given by Borkar in [5][Chap 2].

4.2.1 Assumptions

We denote by \mathcal{F}_n the σ -algebra generated by $(x_0, M_0, \dots, x_n, M_n)$. Namely \mathcal{F}_n contains all the information about the history of the algorithm up to time n .

We introduce the following assumptions:

- (A1) (Lipshitz continuity of h) There exists $L \geq 0$ such that for all $x, y \in \mathbb{R}^d$ $\|h(x) - h(y)\| \leq L\|x - y\|$.
- (A2) (Diminishing step sizes) We have that $\sum_{n \geq 0} \epsilon_n = \infty$ and $\sum_{n \geq 0} \epsilon_n^2 < \infty$.
- (A3) (Martingale difference noise) There exists $K \geq 0$ such that for all n we have that $\mathbb{E}[M_{n+1}|\mathcal{F}_n] = 0$ and $\mathbb{E}[|M_{n+1}|^2|\mathcal{F}_n] \leq K(1 + \|x_n\|)$.
- (A4) (Boundedness of the iterates) We have that $\sup_{n \geq 0} \|x_n\| < \infty$ almost surely.
- (A5) (Liapunov function) There exists a positive, radially unbounded, continuously differentiable function $V : \mathbb{R}^d \rightarrow \mathbb{R}$ such that for all $x \in \mathbb{R}^d$, $\langle \nabla V(x), h(x) \rangle \leq 0$ with strict inequality if $V(x) \neq 0$.

(A1) is necessary to ensure that the o.d.e. has a unique solution given an initial condition, and that the value of the solution after a given amount of time depends continuously on the initial condition. (A2) is necessary for almost sure convergence, and holds in particular for $\epsilon_n = 1/n$. (A3) is required to control the random fluctuations of x_n around the solution of the o.d.e. (using the martingale convergence theorem), and holds in particular if $\{M_n\}_{n \in \mathbb{N}}$ is independent with bounded variance. (A4) is essential, and can (in some cases) be difficult to prove. We will discuss how to ensure that (A4) holds in the latter sections. (A5) ensures that all solutions of the o.d.e. converge to the set of zeros of V , and that this set is stable (in the sense of Liapunov). Barely assuming that all solutions of the o.d.e. converge to a single point does not guarantee convergence of the corresponding stochastic approximation.

4.2.2 The main theorem

We are now equipped to state the main theorem.

Theorem 4.2.1 *Assume that (A1) - (A5) hold, then we have that:*

$$V(x_n) \rightarrow_{n \rightarrow \infty} 0, \text{ a.s.}$$

The proof of theorem 4.2.1 is based on an intermediate result stating that the sequence $\{x_n\}$ (suitably interpolated) remains arbitrarily close to the solution of the o.d.e. We define $\Phi_t(x)$ the value at t of the unique solution to the o.d.e. starting at x at time 0. Φ is uniquely defined because of (A1) and the Picard-Lindelöf theorem. We define $t(n) = \sum_{k=0}^{n-1} \epsilon_k$, and $\bar{x}(t)$ the interpolated version of $\{x_n\}_{n \in \mathbb{N}}$. Namely for all n , $\bar{x}(t(n)) = x_n$, and \bar{x} is linear by parts. We define $x^n(t) = \Phi_{t-t(n)}(x_n)$ the o.d.e. trajectory started at x_n at time $t(n)$.

Lemma 4.2.2 *For all $T > 0$, we have that:*

$$\sup_{t \in [t(n), t(n)+T]} \|\bar{x}(t) - x^n(t)\| \rightarrow_{n \rightarrow \infty} 0 \text{ a.s.}$$

Proof. [Proof of lemma 4.2.2:]

Since the result holds almost surely we consider a fixed sample path throughout the proof. Define $m = \inf\{k : t(k) > t(n) + T\}$ so that we can prove the result for $T = t(m) - t(n)$ and consider the time interval $[t(n), t(m)]$. Consider $n \leq k \leq m$, we are going to start by bounding the difference between \bar{x} and x^n at time instants $t \in \{t(n), \dots, t(m)\}$, that is $\sup_{n \leq k \leq m} |x_k - x^n(t(k))|$.

We start by re-writing the definition of x_k and $x^n(t(k))$:

$$x_k = x_n + \sum_{u=n}^{k-1} \epsilon_u h(x_u) + \sum_{u=1}^{k-1} \epsilon_k M_k$$

and by the fundamental theorem of calculus:

$$\begin{aligned}
x^n(t(k)) &= x_n + \int_{t(n)}^{t(k)} h(x^n(v)) dv \\
&= x_n + \sum_{u=n}^{k-1} \int_{t(u)}^{t(u+1)} h(x^n(v)) dv \\
&= x_n + \sum_{u=n}^{k-1} \epsilon_u h(x^n(t(u))) + \int_{t(u)}^{t(u+1)} [h(x^n(v)) - h(x^n(t(u)))] dv
\end{aligned}$$

we recall that $\int_{t(u)}^{t(u+1)} dv = \epsilon_u$.

Therefore:

$$x^n(t(k)) - x_k = \sum_{u=1}^{k-1} \epsilon_k M_k + \sum_{u=n}^{k-1} \epsilon_u (h(x_u) - h(x^n(t(u)))) + \sum_{u=n}^{k-1} \int_{t(u)}^{t(u+1)} [h(x^n(v)) - h(x^n(t(u)))] dv$$

And it is noted that, from the fact that h is Lipschitz continuous, we have the two inequalities:

$$\begin{aligned}
\|h(x_u) - h(x^n(t(u)))\| &\leq L \|x^n(t(u)) - x_u\| \\
\|h(x^n(v)) - h(x^n(t(u)))\| &\leq L \|x^n(v) - x^n(t(u))\|, \quad v \in [t(u), t(u+1)].
\end{aligned}$$

Our goal is to upper bound the following difference, decomposed into 3 terms:

$$C_k = \|x^n(t(k)) - x_k\| \leq A_k + \sum_{u=n}^{k-1} L B_u + \sum_{u=n}^{k-1} L \epsilon_u C_u, \quad (4.1)$$

with:

$$\begin{aligned}
A_k &= \left\| \sum_{u=n}^{k-1} \epsilon_k M_k \right\|, \\
B_u &= \int_{t(u)}^{t(u+1)} \|x^n(v) - x^n(t(u))\| dv.
\end{aligned}$$

The stochastic term

We first upper bound A_k , the stochastic term in (4.1). Define $S_n = \sum_{u=0}^n \epsilon_u M_u$. It is noted that $A_k = S_k - S_n$. S_n is a martingale since:

$$\mathbb{E}[S_{n+1} - S_n | \mathcal{F}_n] = \mathbb{E}[\epsilon_{n+1} M_{n+1} | \mathcal{F}_n] = 0.$$

From (A3), $\mathbb{E}[|M_{n+1}|^2 | \mathcal{F}_n] \leq K(1 + \sup_k \|x_k\|) < \infty$. Therefore the sequence $\{S_n\}$ is a square integrable martingale:

$$\sum_{n \geq 0} \mathbb{E}[|S_{n+1} - S_n|^2 | \mathcal{F}_n] \leq K(1 + \sup_n \|x_n\|) \sum_{n \geq 0} \epsilon_n^2 < \infty.$$

Using the martingale convergence theorem (lemma 4.3.3), we have that S_n converges almost surely to a finite value S_∞ . This implies that:

$$A_k \leq \|S_k - S_n\| \leq \|S_k - S_\infty\| + \|S_n - S_\infty\| \leq 2 \sup_{n' \geq n} \|S_{n'} - S_\infty\| \xrightarrow{n \rightarrow \infty} 0, \text{ a.s.}$$

Therefore, until the end of the proof we choose n large enough so that $A_k \leq \delta/2$ for all $k \geq n$ with $\delta > 0$ arbitrarily small.

The discretization term, maximal slope of x^n

In order to upper bound B_u , we prove that for $t \in [t(u), t(u+1)]$, $x^n(t)$ can be approximated by $x^n(t(u))$ (up to a term proportional to ϵ_u). To do so we have to bound the maximal slope of $t \mapsto x^n(t)$ on $[t(n), t(m)]$. We know that $x^n(t(n)) = x_n \leq \sup_{n \in \mathbb{N}} \|x_n\|$ which is finite by (A4). Using the fact that h is Lipschitz and applying Gromwall's inequality (lemma 4.3.1) there exists a constant $K_T > 0$ such that:

$$h(x^n(t)) \leq K_T, \quad t \in [t(n), t(m)].$$

We have used the fact that h is Lipschitz so it grows at most linearly: for all x , $\|h(x) - h(0)\| \leq L\|x\|$, so that $\|h(x)\| \leq \|h(0)\| + L\|x\|$. Therefore by the fundamental theorem of calculus, for $t \in [t(u), t(u+1)]$:

$$\|x^n(t) - x^n(t(u))\| \leq \int_{t(u)}^{t(u+1)} \|h(x^n(v))\| dv \leq \epsilon_u K_T.$$

In turn, integrating over $[t(u), t(u+1)]$:

$$B_u \leq \int_{t(u)}^{t(u+1)} \|x^n(v) - x^n(t(u))\| dv \leq \epsilon_u^2 K_T.$$

By (A2), $\sum_{u \geq n} \epsilon_u^2 \xrightarrow{n \rightarrow +\infty} 0$, and so $\sum_{u=n}^{k-1} B_u \leq \sum_{u \geq n} B_u \xrightarrow{n \rightarrow +\infty} 0$. Until the end of the proof, we consider n large enough so that $\sum_{u \geq n} \bar{L} B_u \leq \delta/2$.

The recursive term

Going back to (4.1), by the reasoning above, we have proven that:

$$C_k \leq \delta + L \sum_{u=0}^{k-1} \epsilon_u C_u.$$

Using the fact that $\sum_{u=n}^{k-1} \epsilon_u \leq t(m) - t(n) = T$, and applying the discrete time version of Gronwall's inequality (lemma 4.3.2):

$$\sup_{n \leq k \leq m} C_k \leq \delta e^{LT}.$$

By letting δ arbitrary small we have proven that:

$$\sup_{n \leq k \leq m} \|x_k - x^n(t(k))\| \rightarrow_{n \rightarrow \infty} 0$$

Error due to linear interpolation

In order to finish the proof, we need to provide an upper bound for $\|\bar{x}(t) - x^n(t)\|$ when $t \notin \{t(n), \dots, t(m)\}$. Consider $n \leq k \leq m$, and $t \in [t(k), t(k+1)]$. Since \bar{x} is linear by parts (by definition), there exists $\lambda \in [0, 1]$ such that:

$$\bar{x}(t) = \lambda x_k + (1 - \lambda)x_{k+1}.$$

Applying the fundamental theorem of calculus twice, $x^n(t)$ can be written:

$$\begin{aligned} x^n(t) &= x^n(t(k)) + \int_{t(k)}^t h(x^n(v))dv \\ &= x^n(t(k+1)) - \int_t^{t(k+1)} h(x^n(v))dv \end{aligned}$$

Therefore the error due to linear interpolation can be upper bounded as follows:

$$\begin{aligned} \|\bar{x}(t) - x^n(t)\| &\leq \lambda \|x_k - x^n(t(k))\| + (1 - \lambda) \|x_{k+1} - x^n(t(k+1))\| \\ &\quad + \lambda \int_{t(k)}^t \|h(x^n(v))\|dv + (1 - \lambda) \int_t^{t(k+1)} \|h(x^n(v))\|dv, \end{aligned}$$

and we obtain the announced result:

$$\sup_{t \in [t(n), t(m)]} \|\bar{x}(t) - x^n(t)\| \leq \sup_{n \leq k \leq m} \|x_k - x^n(t(k))\| + \epsilon_n C_T \rightarrow_{n \rightarrow +\infty} 0$$

which concludes the proof. □

We can proceed to prove the main theorem.

Proof.

Once again we work with a fixed sample path. We consider $\nu > 0$, and define the level set $H^\nu = \{x : V(x) \geq \nu\}$. Choose $\epsilon > 0$ such that if $V(x) \leq \nu$ and $\|x - y\| \leq \epsilon$, then $V(y) \leq 2\nu$. Such an ϵ exists because (by radial unboundedness) the set $\{x : V(x) \leq \nu\}$ is compact, and because of the uniform continuity of

V on compact sets. Since V is continuously differentiable, $x \mapsto \langle \nabla V, h(x) \rangle$ is strictly negative on H^ν , and H^ν is closed, we define $\Delta = \sup_{x \in H^\nu} \langle \nabla V, h(x) \rangle < 0$. Denote by $V_\infty = \sup_{\|x\| \leq \sup_n \|x_n\|} V(x)$ which is finite since $\sup_n \|x_n\|$ is finite and V is continuous. Define $T = (V_\infty - \nu)/\Delta$. Then for all x such that $\|x\| \leq \sup_n \|x_n\|$ and all $t > T$, we must have $V(\Phi_t(x)) \leq \nu$. Finally, choose n large enough so that $\sup_{t \in [t(n), t(n)+T]} \|\bar{x}(t) - x^n(t)\| \leq \epsilon$ and m such that $t(m) = t(n) + T$. Then we have that $V(x^n(t+T)) \leq \nu$ and that $|x^n(t+T) - x_m| \leq \epsilon$, which proves that $V(x_m) \leq 2\nu$.

The reasoning above holds for all sample paths, for all $\nu > 0$, for all m arbitrarily large, so $V(x_n) \rightarrow_{n \rightarrow \infty} 0$ a.s. which is the announced result. \square

4.3 Intermediate results

4.3.1 Ordinary differential equations

We state here two basic results on o.d.e.'s used in the proof of the main theorem.

Lemma 4.3.1 (Gronwall's inequality) *Consider $T \geq 0$, $L \geq 0$ and a function $t \mapsto x(t)$ such that $\dot{x}(t) \leq Lx(t)$, $t \in [0, T]$. Then we have that, for all $t \in [0, T]$, $x(t) \leq x(0)e^{Lt}$.*

Proof. Define the function $y(t) = x(t) \exp(-Lt)$. Differentiating we obtain:

$$\dot{y}(t) = (\dot{x}(t) - x(t)L) \exp(-Lt) \leq 0.$$

Hence $t \mapsto y(t)$ is decreasing so that:

$$x(0) = y(0) \geq y(t) = x(t) \exp(-Lt),$$

and for all $t \in [0, T]$ we have $x(t) \leq x(0) \exp(Lt)$ which is the announced result. \square

Lemma 4.3.2 (Gronwall's inequality, discrete case) *Consider $K \geq 0$ and positive sequences $\{x_n\}$, $\{\epsilon_n\}$ such that for all $0 \leq n \leq N$:*

$$x_{n+1} \leq K + \sum_{u=0}^n \epsilon_u x_u.$$

Then we have the upper bound: $x_n \leq K \exp(\sum_{u=0}^{n-1} \epsilon_u)$, for all $0 \leq n \leq N$.

Proof. We are going to prove the stronger result (for all $0 \leq n \leq N$):

$$x_n \leq K \prod_{u=0}^{n-1} (1 + \epsilon_u) \leq K \exp\left(\sum_{u=0}^{n-1} \epsilon_u\right). \quad (4.2)$$

using the elementary inequality $1 + x \leq e^x$. The inequality (4.2) holds for $n = 0$. Let us assume that it holds up to n . Then:

$$\begin{aligned}
x_{n+1} &\leq K + \sum_{u=0}^n \epsilon_u x_u \\
&\leq K \left(1 + \sum_{u=0}^n \epsilon_u \prod_{v=0}^{u-1} (1 + \epsilon_v) \right) \\
&\leq K \left(\sum_{u=0}^n \prod_{v=0}^{u-1} (1 + \epsilon_v) + \sum_{u=0}^n \epsilon_u \prod_{v=0}^{u-1} (1 + \epsilon_v) \right) \\
&\leq K \prod_{u=0}^n (1 + \epsilon_u).
\end{aligned}$$

which proves the result. □

4.3.2 Martingales

We state the martingale convergence which is required to control the random fluctuations of the stochastic approximation in the proof of the main theorem.

Consider a sequence of σ -fields $\mathcal{F} = (\mathcal{F}_n)_{n \in \mathbb{N}}$, and $\{M_n\}_{n \in \mathbb{N}}$ a sequence of random variables in \mathbb{R}^d . We say that $\{M_n\}_{n \in \mathbb{N}}$ is a \mathcal{F} -martingale if M_n is \mathcal{F}_n -measurable and $\mathbb{E}[M_{n+1} | \mathcal{F}_n] = M_n$. The following theorem (due to Doob) states that if the sum of squared increments of a martingale is finite (in expectation), then this martingale has a finite limit a.s.

Theorem 4.3.3 (Martingale convergence theorem) *Consider $\{M_n\}_{n \in \mathbb{N}}$ a martingale in \mathbb{R}^d with:*

$$\sum_{n \geq 0} \mathbb{E}[\|M_{n+1} - M_n\|^2 | \mathcal{F}_n] < \infty,$$

then there exists a random variable $M_\infty \in \mathbb{R}^d$ such that $\|M_\infty\| < \infty$ a.s. and $M_n \xrightarrow{n \rightarrow \infty} M_\infty$ a.s.

Part II

Analysis of wireless networks

Chapter 5

The ALOHA protocol

In this chapter we start our study of distributed multiple access protocols with the simplest of them: the ALOHA protocol. ALOHA (which is a greeting in the Hawaiian language) was developed in the 1970s by N.Abramson ([1]) and others at university of Hawaii. The goal was to design a simple, distributed, low-cost network to connect terminals placed on various islands to a central computer, through a shared wireless channel, in the UHF frequency band (300 MHz - 3 GHz).

The principle of ALOHA could be summarized as follows: when multiple transmitters must share the wireless medium, any transmitter with data to send sends a packet with a fixed probability (which is the protocol parameter), and if a collision is detected, the packet is kept in a buffer for later retransmission. Despite its simplicity, the advantage of such a protocol is that it allows to transmit data in a fully distributed manner, without even being aware of the potential number of competing transmitters. The only requirement in terms of feedback to the transmitter is whether or not a collision has occurred (either through sensing or ACK/NACK).

Although ALOHA is primitive and mostly of historical interest, the analysis of its stability is surprisingly challenging and has been the subject of active research since the original contribution by Tsybakov and Mikailov [31] in 1979. The reason for this difficulty is that, to analyse ALOHA (and more sophisticated multi-access protocols such as CSMA), one must consider queuing systems with multiple servers whose service rate are tightly coupled. Those systems are mostly non-reversible and intractable so that deriving the stationary distribution is not feasible, even for two users and Bernoulli arrivals.

5.1 Packet multiple access

ALOHA is a protocol designed for a particular setting in wireless communication called packet multiple access. The setting is the following: we consider N transmitters trying

to communicate with a single receiver through a shared wireless medium. Time is slotted, and each transmitter transmits packets whose duration is exactly one time slot. For each time slot, three possible events can occur:

- (idle medium) No transmitter attempts to transmit a packet.
- (successful transmission) A *single* transmitter attempts to send a packet. The packet is successfully received and the transmitter is informed of the successful transmission.
- (packet collision) At least two transmitters attempt to send a packet. All the sent packets are lost due to interference at the receiver, and the receiver informs all the transmitters that there has been a collision.

Each transmitter has a buffer where packets to be sent are kept. The buffer contains both packets that have never been transmitted and packets for which a transmission was attempted and a collision occurred so that those packets require retransmission.

Consider a given transmitter, her decision to transmit or to stay silent can be a function of previous collisions/successes and of the current number of packets in her buffer, but is not allowed to depend on anything else. In particular her decisions cannot depend on the number of concurrent transmitters, their past decisions or the state of their buffer. In the context of 802.11-type networks, packet multiple access naturally occurs when several stations try to communicate with a single access point.

ALOHA is a simple protocol designed for packet multiple access, and operates as follows:

- At time slot n , transmitter i draws Y_n^i , a Bernoulli variable with parameter p_i , which is independent of the other users decisions, and the events that have occurred prior to time slot n .
- If $Y_n^i = 1$, a packet is sent, otherwise no packet is sent.
- If a collision is detected, the sent packet is kept in the buffer, and otherwise it is removed from the buffer.

5.2 ALOHA, i.i.d. Bernoulli model

In this section we introduce a simple mathematical model to study the performance of ALOHA, which will be used in the latter sections. We consider N users indexed by $i \in \{1, \dots, N\}$. We define $p_i \in [0, 1]$ the probability for user i to transmit, and λ_i the arrival rate of packets to the buffer of transmitter i . At time slot n , for transmitter i , we define X_n^i the number of packets currently in her buffer, $Z_n^i \in \{0, 1\}$ the number of packets she successfully transmitted, and $A_n^i \in \{0, 1\}$ the number of packets that arrived to her buffer. Define the random variable:

$$Z_n^i = \mathbf{1}\{X_n^i > 0\} Y_n^i \prod_{j \neq i} (1 - Y_n^j \mathbf{1}\{X_n^j > 0\}),$$

so that $Z_n^i = 1$ iff transmitter i has successfully transmitted a packet at time n . It should be noted that transmitter i effectively transmits a packet iff $Y_n^i \mathbf{1}\{X_n^i > 0\} = 1$. The buffers evolve according to the following recursive equations (for all $1 \leq i \leq N$ and $n \in \mathbb{N}$):

$$X_{n+1}^i = X_n^i + A_n^i - Z_n^i.$$

We make the following statistical assumptions on the arrivals and transmissions/collisions:

- For all i , $(A_n^i)_n$ is i.i.d. Bernoulli with parameter λ_i
- $(A_n^i)_n$ is independent of $(A_n^j)_n$ for all $i \neq j$
- For all i , $(Y_n^i)_n$ is i.i.d. Bernoulli with parameter p_i
- $(Y_n^i)_n$ is independent of $(Y_n^j)_n$ for all $i \neq j$

Exercise:

1. Prove that $(X_n)_n$ is a Markov chain on \mathbb{N}^N and calculate its transition probabilities
2. Prove that even for $N = 2$, $(X_n)_n$ is not reversible. Hint: consider the graph associated to the transition matrix of $(X_n)_n$.

5.3 Full buffer analysis

We start by analysing the full-buffer case, where all transmitters have data to transmit at each time slot. This can be done simply by setting $X_0^i = 1$ and $\lambda_i = 1$ so that a packet arrives at each time slot a.s. Namely $A_n^i = 1$ a.s. so that $X_n^i > 0$ a.s. The full-buffer analysis will serve as a preliminary to our the stability analysis, in order to understand how the system behaves in overload. It is understood that the system cannot be stable in this regime, so that we will not be interested in the buffer content $(X_n)_n$, but rather in the throughput T_i (expected number of packets transmitted per unit of time) and the delay (the number of time slots between two successful transmissions).

Both the throughput and delay can be found by simple calculations and are given below.

Proposition 16 (i) *The throughput of user i is $T_i = p_i(1 - \prod_{j \neq i} p_j)$*

(ii) *The delay between two successful transmissions is geometrically distributed with parameter T_i .*

(iii) *Define $T = \sum_{i=1}^N T_i$ the total throughput. Consider the homogeneous case $p_i = p$ for all i . Then the optimal transmission probability $p^* = \arg \max_p T(p)$ is $p^* = 1/N$ and $T(p^*) \xrightarrow{N \rightarrow \infty} e^{-1}$.*

Proof. (i) We consider the full buffer case, so that:

$$T_i = \mathbb{E}[Z_n^i] = \mathbb{E}[Y_n^i(1 - \prod_{j \neq i} Y_n^j)] = p_i(1 - \prod_{j \neq i} p_j).$$

(ii) The delay to transmit the first packet is $\inf\{n \geq 0 : Z_n^i = 1\}$. In the full buffer case, random variables $(Z_n^i)_n$ are i.i.d Bernoulli distributed with parameter T_i so the packet delay is geometrically distributed with parameter T_i , by definition of the geometrical distribution.

(iii) In the homogeneous case, the total throughput is $T(p) = Np(1-p)^{N-1}$. Since $T(0) = T(1) = 0$ and that $T(p) > 0$ for $0 < p < 1$ we must have $\left. \frac{\partial T(p)}{\partial p} \right|_{p=p^*} = 0$ so that:

$$(1-p^*)^{N-1} - p^*(N-1)(1-p^*)^{N-2} = 0.$$

Dividing both sides by $(1-p^*)^{N-2}$ we obtain $(1-p^*) = p^*(N-1)$ so that $p^* = 1/N$. The throughput is:

$$T(p^*) = (1-1/N)^{N-1} = e^{(N-1)\log(1-1/N)} \xrightarrow{N \rightarrow \infty} e^{-1}$$

□

The above result states that the total throughput of ALOHA for a large number of transmitters tends to e^{-1} . We further see that to achieve the optimal performance, the transmission probabilities should be small and inversely proportional to the number of transmitters. Since $e^{-1} \approx 0.37$, we can see that ALOHA is approximately 2.7 times less efficient than time division, where each transmitter is allowed to transmit during a fraction $1/N$ of the time. Indeed time division would yield a total throughput of 1.

5.4 Stability Region of ALOHA

We now turn to the more challenging problem determining the stability region. The stability region of ALOHA is the set of arrival rate vectors $(\lambda_1, \dots, \lambda_N)$ such that there exists a transmission probability vector $p = (p_1, \dots, p_n)$ ensuring that the Markov chain $(X_n)_n$ is positive recurrent, i.e the corresponding queuing system is stable.

The problem of the stability region of ALOHA was first studied in [31]. For the sake of simplicity we state the stability region for Bernoulli i.i.d arrivals, and two transmitters. In this case there exists a simple argument based on stochastic dominance. Namely, one considers an alternative system, where transmitter 2 always transmits with probability p_2 (when her buffer is empty she transmits dummy packets), and study the stability of this alternative system. We then prove that stability of this alternative system implies stability of the original system.

For two users, the original system is described by the recursive equations:

$$\begin{aligned} X_{n+1}^i &= X_n^i + A_n^i - Z_n^i \\ Z_n^1 &= \mathbf{1}\{X_n^1 > 0\}Y_n^1(1 - Y_n^2\mathbf{1}\{X_n^2 > 0\}) \\ Z_n^2 &= \mathbf{1}\{X_n^2 > 0\}Y_n^2(1 - Y_n^1\mathbf{1}\{X_n^1 > 0\}). \end{aligned}$$

The stability region is given by the next result. The proof technique originates in [25]. We define:

$$\rho_i = \frac{\lambda_i}{p_i \prod_{j \neq i} (1 - p_j)},$$

which is the load of transmitter i (the ratio between arrival rate and service rate) in an alternate system where any transmitter $j \neq i$ transmits with probability p_j even when her queue is empty. We recall that the load of a queue is defined as the ratio between the arrival rate and the service rate.

Proposition 17 Consider $N = 2$ transmitters.

(i) The Markov chain $(X_n)_n$ is positive recurrent if either:

$$\lambda_1 < p_1(1 - p_2) \quad \text{and} \quad \lambda_2 < p_2(1 - \rho_1 p_1)$$

or:

$$\lambda_2 < p_2(1 - p_1) \quad \text{and} \quad \lambda_1 < p_1(1 - \rho_2 p_2).$$

(ii) The stability region is given by:

$$\Lambda = \{(\lambda_1, \lambda_2) \in (\mathbb{R}^+)^2 : \sqrt{\lambda_1} + \sqrt{\lambda_2} < 1\}.$$

Proof. (i) The proof is based on a stochastic dominance argument. Define the stochastic process $(\bar{X}_n)_n$:

$$\begin{aligned} \bar{X}_{n+1}^i &= \bar{X}_n^i + A_n^i - \bar{Z}_n^i \\ \bar{Z}_n^1 &= \mathbf{1}\{\bar{X}_n^1 > 0\} Y_n^1 (1 - Y_n^2) \\ \bar{Z}_n^2 &= \mathbf{1}\{\bar{X}_n^2 > 0\} Y_n^2 (1 - Y_n^1 \mathbf{1}\{\bar{X}_n^1 > 0\}). \end{aligned}$$

Essentially, $(\bar{X}_n)_n$ describes a system in which transmitter 2 always transmits with probability p_2 , even when her buffer is empty. When her buffer is empty, she transmits dummy packets. As an intermediate result, we prove that $(\bar{X}_n)_n$ dominates $(X_n)_n$ on each sample path, as stated by the next lemma.

Lemma 5.4.1 For all n : $(X_n^1, X_n^2) \leq (\bar{X}_n^1, \bar{X}_n^2)$ a.s.

Proof. Define $V_n^i = \bar{X}_n^i - X_n^i$. We have $V_{n+1}^i - V_n^i = Z_n^i - \bar{Z}_n^i$.

We first focus on the first transmitter $i = 1$. By definition:

$$\begin{aligned} \bar{Z}_n^1 &= \mathbf{1}\{\bar{X}_n^1 > 0\} Y_n^1 (1 - Y_n^2) \\ Z_n^1 &= \mathbf{1}\{X_n^1 > 0\} Y_n^2 (1 - Y_n^1 \mathbf{1}\{X_n^2 > 0\}). \end{aligned}$$

We have that $\mathbf{1}\{X_n^2 > 0\} \leq 1$ so that:

$$\begin{aligned} 1 - Y_n^2 &\leq 1 - Y_n^2 \mathbf{1}\{X_n^2 > 0\}, \\ Y_n^1(1 - Y_n^2) &\leq Y_n^1(1 - Y_n^2 \mathbf{1}\{X_n^2 > 0\}), \end{aligned}$$

Consider n such that $V_n^1 = 0$ then $\bar{X}_n^1 = X_n^1$, so that, by the above inequality $\bar{Z}_n^1 \leq Z_n^1$. This proves that if $V_n^1 = 0$ then $V_{n+1}^1 = V_n^1 + Z_n^1 - \bar{Z}_n^1 \geq V_n^1 = 0$. Since $V_0^1 = 0$, we have proven by induction that $V_n^1 \geq 0$ for all n , so that $\bar{X}_n^1 \geq X_n^1$ for all n .

Let us consider the second transmitter $i = 2$. We have proven that for all n : $X_n^1 \leq \bar{X}_n^1$ so that $\mathbf{1}\{X_n^1 > 0\} \leq \mathbf{1}\{\bar{X}_n^1 > 0\}$ and hence:

$$Y_n^2(1 - Y_n^1 \mathbf{1}\{\bar{X}_n^1 > 0\}) \leq Y_n^2(1 - Y_n^1 \mathbf{1}\{X_n^1 > 0\}).$$

By the same reasoning as above, $V_n^2 = 0$ implies that $\bar{Z}_n^2 \leq Z_n^2$ and $V_{n+1}^2 = V_n^2 + Z_n^2 - \bar{Z}_n^2 \geq V_n^2 = 0$. Since $V_0^2 = 0$, by induction $V_n^2 \geq 0$ for all n , so that $\bar{X}_n^2 \geq X_n^2$ for all n .

Finally, have proven that for all n : $(X_n^1, X_n^2) \leq (\bar{X}_n^1, \bar{X}_n^2)$, a.s. which is the announced result. \square

First consider the stochastic process $(\bar{X}_n^1)_n$. From its definition one sees that $(\bar{X}_n^1)_n$ is independent of $(\bar{X}_n^2)_n$. Furthermore, $(\bar{X}_n^1)_n$ is a discrete M/M/1 queue with arrival rate λ_1 and service rate $p_1(1 - p_2)$. Its load is $\rho_1 = \frac{\lambda_1}{p_1(1-p_2)}$ as defined above. Therefore $(\bar{X}_n^1)_n$ is positive recurrent iff $\rho_1 < 1$, that is:

$$\lambda_1 < p_1(1 - p_2)$$

Furthermore we have that under the above condition, $(\bar{X}_n^1)_n$ is an ergodic Markov chain, with $\mathbb{P}[\bar{X}_n^1 > 0] = \rho_1$.

Now turn to stochastic process $(\bar{X}_n^2)_n$. $(\bar{X}_n^2)_n$ is a one-dimensional queue with arrival rate λ_2 and time-varying service rate. Its expected service rate is given by:

$$p_2(1 - p_1 \mathbb{P}[\bar{X}_n^1 > 0]) = p_2(1 - \rho_1 p_1)$$

Using Loynes' theorem for queues with stationary ergodic arrivals and service times [2], $(\bar{X}_n^2)_n$ is positive recurrent iff¹:

$$\lambda_2 < p_2(1 - \rho_1 p_1).$$

¹ The stability of a general stationary ergodic queue is beyond the scope of these lecture notes, so that Loynes's theorem will be used without proof. The interested reader might consult [2] which covers the topic in details.

We have previously proven that $(X_n^1, X_n^2) \leq (\bar{X}_n^1, \bar{X}_n^2)$ a.s. , so that positive recurrence of $(\bar{X}_n)_n$ implies positive recurrence of $(X_n)_n$, and we have seen that a sufficient condition is:

$$\lambda_1 < p_1(1 - p_2) \quad , \quad \lambda_2 < p_2(1 - \rho_1 p_1).$$

By interchanging the roles of 1 and 2, we obtain another sufficient condition:

$$\lambda_2 < p_2(1 - p_1) \quad , \quad \lambda_1 < p_1(1 - \rho_2 p_2).$$

(ii) The second statement is obtained by considering $\lambda_1 \in [0, 1)$ fixed and deriving the maximal allowed value of λ_2 by varying parameters (p_1, p_2) . (left as an exercise) \square

The proof technique suggests the following remark. In the alternate system, one of the queues (say $(X_n^1)_n$) is independent of the other $((X_n^2)_n)$. Hence, seen from $(X_n^2)_n$, $(X_n^1)_n$ is a stationary process which is exogenous and determines the (time varying) probability that 2 successfully transmits a packet. This is in fact a particular case of a mean field argument. A generic mean field argument would be to consider user i , and assume that, from the point of view of i , any other queue $(X_n^j)_n$ (with $j \neq i$) is a stationary process and independent of $(X_n^i)_n$. Surprisingly enough, a mean field argument can be used to yield bounds on the stability region of ALOHA [4], and those bounds become tight when $N \rightarrow \infty$, so that when the number of users is large, we may retrieve the stability region exactly.

It is also interesting to compare the stability region of ALOHA to that of time division:

$$\begin{aligned} \Lambda^{ALOHA} &= \{(\lambda_1, \lambda_2) \in (\mathbb{R}^+)^2 : \sqrt{\lambda_1} + \sqrt{\lambda_2} < 1\} \\ \Lambda^{TDMA} &= \{(\lambda_1, \lambda_2) \in (\mathbb{R}^+)^2 : \lambda_1 + \lambda_2 < 1\}. \end{aligned}$$

Chapter 6

The CSMA protocol

We now focus on CSMA as it is implemented in 802.11-type networks. Due to the presence of collisions and the strong coupling between the various transmitters calculating the performance of CSMA is most likely intractable. We give an exposition of Bianchi's approximate analysis, which allows to calculate the performance of CSMA in closed-form with a very small approximation error. To understand the engineering of CSMA, we study some regimes of practical interest where the throughput and optimal window size have a simple expression.

6.1 The CSMA algorithm

6.1.1 CSMA principles

Consider once again the packet multiple access problem, which we first studied in the analysis of ALOHA. The reason behind the poor performance of ALOHA when compared to TDMA (Time Division Multiple Access) is the fact that decisions at successive instants are not correlated, in other words ALOHA does not make use of the feedback available after each transmission attempt. Two types of feedback are available to a transmitter: (i) the transmitter can hear other transmitters in her vicinity (ii) after each transmission attempt, a transmitter is informed on whether or not the sent packet was successfully decoded (by ACK, NACK or and absence of ACK/NACK).

Carrier Sense Multiple Access (CSMA) is a family of algorithms which makes use of this feedback information. In fact the basic principles of CSMA closely follow the social conventions used by (polite) adults when attempting to have a conversation with a large number of participants, for instance in a classroom or at a party. The principles can be phrased as follows:

- (1) Listen before talking: Each transmitter monitors the medium, and does not transmit when the medium is sensed busy i.e when she hears another transmission

taking place

- (2) Give everyone a chance to speak: When a transmitter is done transmitting a packet, she waits for the medium to appear idle for a sufficiently long amount of time. This principle applies even when the buffer of the transmitter is not empty. This guarantees that no transmitter hogs the medium after gaining access to it.
- (3) Raise your hand if you have a question: Before attempting to seize the medium, a transmitter warns the receiver by transmitting a short packet called Request To Send (RTS). The receiver, in turn, broadcasts this message to all concurrent transmitters that are within her range by transmitting a Clear To Send (CTS) packet. This solves the hidden node issue, where transmitters may interfere without being able to hear each other.
- (4) If you interrupt, stop talking and wait: If a transmitter attempts to transmit and detects a collision, she stops transmitting, and waits for a certain amount of time. The amount of time to wait is called the "back-off", and is typically a random variable.

6.1.2 Variants

There exists several variants of CSMA, and the difference between those variants mainly lies in the back-off mechanism and the way that transmitters are warned about incoming transmissions. The main variants are:

- CSMA with Collision Avoidance (CSMA/CA): Principles (1),(2),(3) and (4) are implemented. It should be noted that once a station has seized the medium (Principle (3)), she sends her packet in its entirety, even if a collision is detected between the start and the end of the packet transmission.
- CSMA with Collision Detection (CSMA/CD): Principle (3) is not applied. On the other hand, if a transmitting station detects a collision while transmitting a packet, she signals it to the other stations by immediately sending a jamming signal, then stops transmitting. Hence CSMA/CD does not solve the hidden node problem.

CSMA/CA is used in WiFi networks (IEEE 802.11), while CSMA/CD is used in Ethernet networks (IEEE 802.03). This choice is natural since CSMA/CA is designed to overcome the hidden node problem which arises only in wireless scenarios. In wired networks, all transmitters are connected by a hub (a repeater), so that two transmitters associated to the same receiver can always hear each other. Also note that, in 802.11 transmitters have the option of not using RTS/CTS frames. In that case one calls this access scheme the "Basic Access Scheme".

6.1.3 Physical and Virtual Sensing

As mentioned before, the medium is sensed by two techniques which complement each other:

- Physical sensing: each transmitter calculates the received power at each instant. If the power level is above a given threshold, the medium is sensed busy, otherwise the medium is sensed idle. It should be noted that physical sensing consumes power.
- Virtual Sensing: when a transmitter transmits a packet, she includes an information on how much time she will be occupying the medium. This information is called NAV (Network Allocation Vector). Any station that hears such a packet does not transmit nor senses the medium until a time equal to NAV has elapsed.

It should be noted that RTS/CTS packets also include the NAV, which means that any station in range of *either* the transmitter or the receiver will be able to hear this information. This also means that, when the RTS/CTS mechanism is employed, collisions may only occur during the transmission of the RTS/CTS packets.

6.1.4 Formal description

Let us now give a formal description of CSMA. Consider N transmitters associated to a single receiver and all sharing the wireless medium. Transmitters are assumed symmetrical, and let us focus on transmitter 1. Consider time indexed by $t \in \mathbb{R}$, and define $b(t) \in \{0, 1\}$ indicating whether or not the medium is sensed busy by transmitter 1. Namely $b(t) = 1$ if transmitter 1 senses the medium to be idle at time t and $b(t) = 0$ otherwise. Define $B = \{t \in \mathbb{R} : b(t) = 0\}$ the set of instants where the medium appears idle to transmitter 1. It should be noted that whenever $t \notin B$, transmitter 1 remains idle: she does not attempt to transmit, and her back-off countdown is frozen.

We are going to map this system into a discrete time system by sampling at the time where certain types of events occur. For the sake of simplicity, we will consider that the duration needed to transmit a packet is fixed. In general the packet size is random, since the payload does not have a fixed size. We recall that the payload is the "useful information" contained in a packet, i.e it is the contents of the packet besides signalling information contained in the headers. We consider 3 types of events:

- (back-off decrement) the transmitter waits for a time slot, and her back-off countdown decreases by 1. This event has duration σ .
- (collision) the transmitter initiates a transmission during which a collision occurs. This event has duration T_c .
- (successful transmission) the transmitter initiates a transmission during which no collision occurs. This event has duration T_s .

Define $\mathcal{N} = \{t_0, t_1, \dots\}$ the (discrete) set of instants at which one of the three events described above start. Now define $(s(n), b(n))$ the state of transmitter 1 at time t_n . $s(n) \in \{1, \dots, m\}$ is called the back-off stage, and $b(n) \in \{1, \dots, W_{s(n)} - 1\}$ is the back-off counter. W_i is the maximal back-off window at back-off stage i . Typically, we will assume that the window size is exponential in the backoff stage so that $W_i = W2^i$. By a slight abuse of terminology, we will refer to W as the “window size” (instead of saying “the window size at back-off stage 0”). Finally, we will define $c(n) \in \{0, 1\}$, so that $c(n) = 0$ denotes the event that, if a packet is transmitted at time t_n , then the transmission will be successful, and $c(n) = 1$ otherwise.

The state of transmitter 1 evolves as follows:

- (back-off decrement) If $b(n) > 0$ then $(s(n+1), b(n+1)) = (s(n), b(n) - 1)$.
- (collision) If $b(n) = 0$ and $c(n) = 1$ then $s(n+1) = \min(s(n) + 1, m)$ and $b(n+1)$ is uniformly distributed in $\{0, \dots, W_{s(n+1)} - 1\}$,
- (collision) If $b(n) = 0$ and $c(n) = 0$ then $s(n+1) = 0$ and $b(n+1)$ is uniformly distributed in $\{0, \dots, W_0 - 1\}$.

6.2 Performance of CSMA: Bianchi’s model

6.2.1 The key assumption

As in the case of ALOHA, we consider the full-buffer regime where each transmitter has a packet to transmit at each time slot. We describe an approximate model proposed by [3] which is both tractable, and surprisingly accurate when compared to numerical experiments. The interested reader might find extensions of Bianchi’s analysis in [19].

By inspecting the description above, we can notice that $(s(n), b(n))_n$ is *not* a Markov chain. Indeed, conditional to $(s(n), b(n))$ the evolution of $s(n+1), b(n+1)$ depends critically on $c(n)$, which is a function of the state of *all* the transmitters. The key idea of Bianchi’s analysis is to assume that $(c(n))_n$ is an i.i.d Bernoulli process, with parameter p (the collision probability). This is of course an approximation of the original model, but it allows to analyse the system using Markov chains. Indeed, if $(c(n))_n$ is i.i.d then the distribution of $s(n+1), b(n+1)$ depends solely on $(s(n), b(n))$, so that $(s(n), b(n))_n$ is a Markov chain.

Once again, if N is large, Bianchi’s idea is close to a mean-field model. Namely the collision process $(c(n))_n$ is the result of interactions between a large number of transmitters, so that, from the point of view of 1, the result of this interaction should “look i.i.d”. Hence Bianchi’s analysis should, at least intuitively be close to the original model for large values of N . Numerical experiments confirm that this is indeed true, and perhaps more surprisingly, the analysis is even quite accurate when N is small.

6.2.2 Stationary distribution

Let us now specify the transition probabilities of the Markov chain, and define $P(s', b'; s, b) = \mathbb{P}[(s(n+1), b(n+1)) = (s', b') | (s(n), b(n)) = (s, b)]$ the transition matrix. By inspection of the description given in 6.1.4 we have:

$$\begin{aligned} P(s, b-1; s, b) &= 1, \quad b > 0 \\ P((s+1) \wedge m, b; s, 0) &= \frac{p}{W_{(s+1) \wedge m}}, \quad b' \in \{0, \dots, W_{(s+1) \wedge m} - 1\} \\ P(0, b'; s, s, 0, 0, b') &= \frac{1-p}{W_0}, \quad b' \in \{0, \dots, W_0 - 1\} \end{aligned}$$

The transition matrix described above is irreducible by inspection, it is positive recurrent since it is finite, and it is aperiodic since $P(m, 0, m, 0) = p/W_m > 0$. It should also be noted that the transition matrix is *not* reversible, since $P(0, 0, 1, 0) = 1$ and $P(0, 0, 0, 1) = 1$.

Proposition 18 *The stationary distribution $\pi = (\pi_{s,b})_{s,b}$ of P is given by:*

$$\pi_{s,b} = p^s (1-p)^{-f_s} (W_s - b) / W_s,$$

with $f_s = \mathbf{1}\{s = m\}$. Furthermore:

$$\pi_{0,0} = \frac{2(1-2p)(1-p)}{W(1-p(1+(2p)^m)) + 1 - 2p}$$

Proof.

Balance equations

Let us now derive the stationary distribution by solving the full balance equations. We proceed by stages, in order to take advantage of the particular structure of the chain. Consider stage $0 < s < m$, then the balance equations give:

$$\begin{aligned} \pi_{s, W_s - 1} &= (p/W_s) \pi_{s-1, 0} \\ \pi_{s,b} &= (p/W_s) \pi_{s-1, 0} + \pi_{s, b+1}, \quad 0 \leq b < W_s - 1. \end{aligned}$$

So that substituting the first equation in the second gives, for $0 \leq b < W_s$:

$$\pi_{s,b} = \pi_{s, W_s - 1} + \pi_{s, b+1}$$

and by induction we deduce that:

$$\pi_{s,b} = (W_s - b) \pi_{s, W_s - 1} = (W_s - b) p \pi_{s-1, 0} / W_s.$$

Setting $b = 0$ we have that $\pi_{s,0} = p\pi_{s-1,0}$ so that by induction $\pi_{s,0} = p^s\pi_{0,0}$.

Let us now consider back-off stage m . We proceed in similar fashion:

$$\pi_{m,W_m-1} = p(\pi_{m-1,0} + \pi_{m,0})/W_m \quad (6.1)$$

$$\pi_{m,b} = p(\pi_{m-1,0} + \pi_{m,0})/W_m + \pi_{m,b+1}, \quad 0 \leq b < W_m - 1 \quad (6.2)$$

Hence for $0 \leq b < W_s$:

$$\pi_{m,b} = \pi_{m,W_m-1} + \pi_{m,b+1},$$

so that by induction $\pi_{m,b} = (W_m - b)\pi_{m,W_m-1}$. Setting $b = 0$ we have that $\pi_{m,0} = W_m\pi_{m,W_m-1}$. It remains to calculate π_{m,W_m-1} . We have proven that $\pi_{m-1,0} = p^{m-1}\pi_{0,0}$, therefore rewriting (6.1) we obtain:

$$\pi_{m,W_m-1} = p(p^{m-1}\pi_{0,0} + W_m\pi_{m,W_m-1})/W_m$$

from which we deduce:

$$\pi_{m,W_m-1} = p^m(1-p)^{-1}\pi_{0,0}$$

Putting it all together, we have proven that:

$$\pi_{s,b} = \pi_{0,0}p^s(W_s - b)/W_s, \quad s < m$$

$$\pi_{m,b} = \pi_{0,0}p^m(1-p)^{-1}(W_m - b)/W_m,$$

which is the announced result.

Normalization constant

We may now calculate $\pi_{0,0}$ by normalization. Indeed:

$$\begin{aligned} \pi_{0,0}^{-1} &= \sum_{s=0}^m p^s(1-p)^{-f_s} \sum_{b=0}^{W_s-1} (W_s - b)/W_s \\ &= \sum_{s=0}^m p^s(1-p)^{-f_s}(W_s + 1)/2 \\ &= \sum_{s=0}^m p^s(1-p)^{-f_s}(W2^s + 1)/2 \end{aligned}$$

Decomposing the r.h.s. into two terms we have:

$$\begin{aligned} \sum_{s=0}^m (2p)^s(1-p)^{-f_s} &= \frac{1 - (2p)^m}{1 - 2p} + \frac{(2p)^m}{1 - p} \\ &= \frac{(1 - (2p)^m)(1 - p) + (2p)^m(1 - 2p)}{(1 - p)(1 - 2p)} \\ &= \frac{1 - p(1 + (2p)^m)}{(1 - p)(1 - 2p)}. \end{aligned}$$

And similarly:

$$\sum_{s=0}^m p^s (1-p)^{-fs} = \frac{1-p^m}{1-p} + \frac{p^m}{1-p} = \frac{1}{1-p}$$

therefore:

$$\pi_{0,0}^{-1} = \frac{W(1-p(1+(2p)^m)) + 1 - 2p}{2(1-2p)(1-p)}$$

which is the announced result. \square

6.2.3 Fixed point equation

There is still a missing element: according to the previous calculation, we are now able to characterize τ , the probability that transmitter 1 attempts to transmit in stationary state as a function of the collision probability p . Indeed:

$$\tau = \mathbb{P}[b(n) = 0] = \sum_{s=0}^m \pi_{s,0} = \pi_{0,0}/(1-p) = \frac{2(1-2p)}{W(1-p(1+(2p)^m)) + 1 - 2p}.$$

But p is not an input parameter, and depends on whether or not the other transmitters attempt to transmit. Since in Bianchi's model, we have assumed that the state of the various transmitters are independent, then p is the probability that at least one of the $N-1$ remaining transmitters attempt to transmit one has:

$$p = 1 - (1 - \tau)^{N-1}.$$

Therefore, to derive the throughput one must solve the following fixed point equation:

$$p = 1 - (1 - \tau)^{N-1}$$

$$\tau = \frac{2(1-2p)}{W(1-p(1+(2p)^m)) + 1 - 2p}.$$

Such an equation can be solved numerically by an iterative scheme. Based on the solution (which is unique), we may now derive the throughput.

6.2.4 Throughput

We may now derive the total throughput of the system. Consider a randomly chosen slot time in stationary state ¹, and recall that we have assumed that all transmitters behave in an independent manner:

¹The reader familiar with point processes and Palm calculus might notice that the sentence "consider a randomly chosen slot time" is ambiguous. What we mean here is "consider the slot time number 0 under the Palm measure".

- The slot is empty with probability $p_e = (1 - \tau)^N$ (duration σ)
- The slot contains a successful transmission with probability $p_s = N\tau(1 - \tau)^{N-1}$ (duration T_c)
- The slot contains a collision with probability $p_c = 1 - (p_s + p_e)$ (duration T_c)

Therefore the expected length of a slot time is $l = p_e\sigma + p_sT_c + p_cT_c$. Whenever there is a successful transmission, the amount of information transferred is the payload of a packet which we will denote by P . Define $N(T)$ the number of time slots between time 0 and T , and $N_s(T)$ the number of time slots where a successful transmission has occurred. From ergodicity we have that $N(T)/T \xrightarrow{T \rightarrow \infty} 1/l$ and $N_s(T)/N(T) \xrightarrow{T \rightarrow \infty} p_s$ a.s. so that the number of successfully transmitted bits over a large time horizon is given by:

$$S = \lim_{T \rightarrow \infty} PN_s(T)/T = \lim_{T \rightarrow \infty} P \frac{N_s(T)}{N(T)} \frac{N(T)}{T} = \frac{p_s P}{p_e\sigma + p_sT_s + p_cT_c}.$$

6.3 Engineering Insights

We may now show how the formulas obtained by Bianchi's model give insight in the basic engineering issues of the CSMA protocol namely:

- What is the achievable throughput for a large number of users ? How far is CSMA from time division ?
- What is the impact of the payload size on the performance ?
- How should the window size and the number of back-off stages be chosen ?
- Does the RTS/CTS mechanism really improve over the basic access scheme ?

6.3.1 Large user regime

We will be concerned with the regime where $N \rightarrow \infty$ is large. In that regime, one must have that the probability of successful transmission $p_s = N\tau(1 - \tau)^{N-1}$ tends to a strictly positive value, so that τ must be of order $1/N$. Consider $\tau = x/N$ with x a value over which we will subsequently optimize. Using the fact that $\lim_{N \rightarrow \infty} (1 + y/N)^N = e^y$, the probabilities of interest become:

$$\begin{aligned} p_e &= (1 - x/N)^N \xrightarrow{N \rightarrow \infty} e^{-x} \\ p_s &= x(1 - x/N)^{N-1} \xrightarrow{N \rightarrow \infty} xe^{-x} \\ p_c &= 1 - (p_e + p_s) \xrightarrow{N \rightarrow \infty} 1 - (1 + x)e^{-x} \\ p &= 1 - (1 - x/N)^{N-1} \xrightarrow{N \rightarrow \infty} 1 - e^{-x} \end{aligned}$$

The throughput tends to:

$$S \xrightarrow{N \rightarrow \infty} \bar{S} = \frac{Pxe^{-x}}{\sigma e^{-x} + xe^{-x}T_s + (1 - (1+x)e^{-x})T_c}.$$

Rewriting the above equation, we get:

$$\bar{S} = P[\sigma/x + T_s + T_c((e^x - 1)/x - 1)]^{-1}$$

Before we attempt to optimize the above expression in x , it is worth noting that the optimal value of x does not depend on T_s , and in fact only depends on the ratio between T_c and σ . We will denote this quantity $T_c^* = T_c/\sigma$. Therefore we aim at minimizing the following quantity:

$$f(x) = \frac{1 + T_c^*(e^x - 1)}{x}.$$

Differentiating f and equating the derivative to 0 gives the following equation for the optimal value x^* :

$$e^{x^*}(x^* - 1) = 1/T_c^* - 1. \quad (6.3)$$

The equation above has a unique solution and can be found numerically using, for instance, Newton's method.

6.3.2 Window size

Let us now derive the corresponding value of the contention window. Recall that:

$$\tau = \frac{2(1 - 2p)}{W(1 - p(1 + (2p)^m)) + 1 - 2p}$$

Hence the value of W is given by:

$$W = \frac{(2/\tau - 1)(1 - 2p)}{1 - p(1 + (2p)^m)}.$$

Recall that $p \xrightarrow{N \rightarrow \infty} e^{-x}$, so that in the asymptotic regime one must have:

$$W \sim_{N \rightarrow \infty} \frac{2N(2e^{-x} - 1)}{x(e^{-x} + 2^m(e^{-x} - 1)^{m+1})}.$$

6.3.3 Small slot regime

In practice, the slot time σ is much smaller than the length of a collision, so that T_c^* is very large. In turn, from (6.3), x^* must be close to 0, and by using a Taylor approximation $e^x = 1 + x + o(x)$ we obtain the simple (approximate) solution:

$$x^* = \sqrt{2/T_c^*}.$$

Hence in that regime one must have $\tau = \sqrt{2/T_c^*}/N$. In turn the optimal window size in the small slot regime is simply:

$$W^* = 2N/x^* = N\sqrt{2T_c^*}.$$

6.4 Typical Parameter Values

Numerically, Bianchi's model predicts a throughput of 0.83. The deviations with simulations are very small 1 – 2%, so that for all practical purposes the throughput performance of CSMA can be calculated in closed form. Furthermore, one may check that the approximate window size $W^* = N\sqrt{2T_c^*}$ gives close to optimal performance.

The following numerical values apply are not meant to be exact but merely to give an idea of the typical values of the system parameters, in order to justify the various regimes studied in the above sections.

Slot size σ	$50\mu s$
Payload P	$8184\mu s$
Transmission time T_s (no RTS/CTS)	$8982\mu s$
Collision time T_c (no RTS/CTS)	$8713\mu s$
Transmission time T_s (with RTS/CTS)	$9568\mu s$
Collision time T_c (with RTS/CTS)	$417\mu s$

Chapter 7

Scheduling

CSMA as discussed in the previous chapter is an example of a *scheduling algorithm*, in the sense that it is a rule that allows to determine which transmitters may use the medium at any given time. In the full buffer regime, neither ALOHA nor CSMA can achieve the maximal throughput (the throughput of time division), due to collisions. Furthermore, when considering queuing dynamics (even when collisions are neglected) both those algorithms are not throughput optimal, namely they do not stabilize the network when it is possible at all. This lack of optimality stems from the fact that the probability that a transmitter obtains access to the channel does not depend on her queue length. In this chapter we shift our focus to the case where collisions do not occur, and discuss algorithms that are throughput optimal. We first treat the Max-Weight algorithm, which is a *centralized*, throughput optimal, scheduling algorithm. We then show that there exists CSMA-like algorithms which achieve throughput optimality in a *fully distributed manner*. These algorithms can be seen as CSMA where the window size of a transmitter is a decreasing function of her queue length. The analysis of their throughput optimality and convergence time is based on stochastic approximation and mixing time of reversible Markov chains.

7.1 Scheduling in constrained queuing systems

7.1.1 Constrained queuing systems

Let us first describe a model for wireless scheduling known as “constrained queuing systems”. We consider N links, where each link is a transmitter-receiver pair. Time is slotted and indexed by n . The number of packets in the buffer of link k (its backlog) at time n is denoted by $Q_k(n)$. We will call $Q(n) = (Q_i(n))_i$ the congestion process. For the sake of simplicity, there is no rate adaptation and all links have rate 1. Namely, if a link transmits during a time slot it transmits exactly one packet.

Due to interference, not all links may be activated at the same time, hence the term “constrained queuing system”. The constraints are represented as follows. We define a schedule $x = (x_1, \dots, x_N) \in \{0, 1\}^N$, where, for all k , $x_k = 1$ iff link k transmits and 0 otherwise. The interference is represented by a graph $G = (V, E)$, with V the links, and $(k, k') \in E$ iff k and k' cannot transmit at the same time i.e. they interfere. A schedule x is feasible iff all active links do not interfere with each other. Namely, $x_k x_{k'} = 1$ implies $(k, k') \notin E$ for all k, k' . Another way of phrasing this is that x must be an independent set of G . Denote by \mathcal{I} the set of all independent sets. \mathcal{I} is finite, denote by $I = |\mathcal{I}|$ its cardinal, and define x^i the i -th element of \mathcal{I} .

We assume that the arrival process of packets to each link is described by a Bernoulli process, and that arrivals at different links are independent. Namely define $A_k(n)$ the number of packets arriving at link k in time slot n , we assume that $(A_k(n))_n$ is a Bernoulli i.i.d. sequence with parameter λ_k . λ_k is the arrival rate of link k .

At each time slot, a different schedule might be chosen, and a policy is a choice of $(x(n))_n$ s.t. $x(n) \in \mathcal{I}$ for all n . Given a policy, the evolution of the backlogs may be written as:

$$Q_i(n+1) = \max(Q_i(n) + A_i(n) - x_i(n), 0).$$

In general, a wireless network cannot be modelled by a constrained queuing system, see discussion above. Constrained queuing systems are acceptable if we are willing to assume the following:

- Interference is represented by the protocol model (which gives the interference graph G)
- All links/nodes have the ability to sense the medium perfectly and instantly
- There are no hidden nodes: for any pair of links $(k, k') \in E$ then either k can hear transmissions from k' or vice-versa.

Hence there are no collisions in such a system, and the main interest of queuing systems, despite being simplistic models of wireless systems, is their *tractability*.

7.1.2 Stability region

In the problem defined above, the main goal is to design *throughput-optimal* policies. A throughput-optimal policy is a policy that stabilizes the network when this is possible at all. In Markov chain terminology, this means that the congestion process $(Q(n))_n$ is positive recurrent. If arrival rates are too large, there exists no policy that stabilizes the network, and the stability region is defined as the set of arrival rates such that there exists a policy guaranteeing positive recurrence of $(Q(n))_n$.

Definition 7.1.1 *The stability region Λ is defined as the closure of the set of λ such that there exists a policy ensuring that $(Q(n))_n$ is positive recurrent.*

In fact the stability region may be calculated using flow-conservation arguments. Indeed, define the I -dimensional simplex:

$$A = \{\alpha = (\alpha_1, \dots, \alpha_I) : \alpha \geq 0, \sum_{i=1}^I \alpha_i \leq 1\}.$$

We denote the dot product by $\langle \dots \rangle$. We further define $X_k = (x_k^i)_i$.

Proposition 19 *The stability region Λ is given by:*

$$\Lambda = \{(\langle \alpha, X_k \rangle)_k : \alpha \in A\}.$$

Proof. Consider λ in the interior of Λ . Then there exists $\lambda < \bar{\lambda} \in A$. Then there exists $\alpha \in A$ such that $\bar{\lambda}_k = \langle \alpha, X_k \rangle$. Consider the policy where $(x(n))_n$ is i.i.d, and $\mathbb{P}[x(n) = x^i] = \alpha_i$. Queue k : arrival rate λ_k and service rate $\langle \alpha, X_k \rangle = \bar{\lambda}_k > \lambda_k$. So queue k is stable. Hence $(Q(n))_n$ is positive recurrent.

Consider λ in the interior of the complement of Λ . Assume that there exists a policy $(x(n))_n$ such that $(Q(n))_n$ is positive recurrent. Then there exists $\alpha \in A$ with $\frac{1}{n} \sum_{n'=1}^n \mathbf{1}\{x(n') = x^i\} \xrightarrow[n \rightarrow \infty]{} \alpha_i$. Queue k is stable so that:

$$\lambda_k \leq \langle \alpha, X_k \rangle,$$

so $\lambda \in \Lambda$, a contradiction. \square

Not only does the above reasoning give stability region, but it also shows that, if λ is known, there exists a throughput optimal policy which is state independent, so that $x(n)$ is independent of $Q(1), \dots, Q(n)$.

7.2 Centralized scheduling: the Max-Weight algorithm

7.2.1 The Max-Weight algorithm

When λ is unknown, there are two reasonable possibilities:

- **Estimation of the arrival rates:** at time n estimate λ by the empirical average $\hat{\lambda}(n) = (1/n) \sum_{n'=1}^n A(n')$, then calculate $\hat{\alpha}(n) \in A$ such that $\hat{\lambda}(n) = \sum_{i=1}^I \hat{\alpha}_i x^i$, and choose $x(n) = x^i$ with probability $\hat{\alpha}_i(n)$. This policy is throughput optimal since $\hat{\lambda}(n) \xrightarrow[n \rightarrow \infty]{} \lambda$ a.s. by the law of large numbers
- **State-dependent policy:** one does not attempt to estimate λ , and chooses schedule based on the *current* backlogs $x(n) = f(Q(n))$ with f well chosen. Typically, to ensure stability one should give priority to links which have a large backlog, while being as greedy as possible when all links have a small backlog.

The first method has two issues: a) it can only work in a stationary environment where arrival rates are unknown but constant. Such an approach would fail when arrival rates are slowly varying, which is typically the case in wireless networks due to higher level protocols such as TCP b) it requires memory, which creates overhead.

An instance of the second method is the celebrated Max-Weight algorithm. The Max Weight algorithm is described by the following equation:

$$x(n+1) = \arg \max_{x \in \mathcal{I}} \langle x, Q(n) \rangle.$$

Several remarks are in order:

- Large queues get priority to limit congestion
- When all queues are equal, the algorithm is greedy: the largest independent set is chosen to maximize the sum of throughputs
- $x(n+1)$ is the maximum weighted independent set of the weighted graph $G = (V, E, Q(n))$, hence the algorithm name.

In fact Max-Weight derives from the following reasoning. From our previous study of the Foster criterion, we know that if $Q(n)$ admits a Liapunov V function then $Q(n)$ is positive recurrent. For a throughput optimal policy, we need the drift of V to be strictly less than 0 for all $\lambda \in \Lambda$. Hence, a reasonable idea is to select a given V , and then, at all times, choose the decision minimizing the drift $\mathbb{E}[V(Q(n+1)) - V(Q(n)) | Q(n)]$. The Max-Weight is in fact an example of this idea with V taken as the square norm of the backlog: $V(Q(n)) = \|Q(n)\|^2$.

7.2.2 Throughput optimality

We now establish throughput optimality of Max-Weight scheduling. As said above, the proof simply involves calculating the drift of the norm of the congestion process $V(Q(n)) = \|Q(n)\|^2$, and then apply Foster's criterion.

Proposition 20 *Max-Weight scheduling is throughput optimal.*

Proof. Consider λ in the interior of Λ . We first calculate the drift:

$$\begin{aligned} \|Q(n+1)\|^2 &= \|Q(n+1) - Q(n) + Q(n)\|^2 \\ &= \|Q(n)\|^2 + 2\langle Q(n), Q(n+1) - Q(n) \rangle + \|Q(n+1) - Q(n)\|^2. \end{aligned}$$

Recall that the backlogs obey the following recursive equation:

$$Q(n+1) = \max(Q(n) + A(n) - x(n+1), 0).$$

We now examine two possible cases:

- If $Q_k(n) + A_k(n) - x_k(n+1) < 0$ then $Q_k(n) = Q_k(n+1) = 0$
- Otherwise: $Q_k(n+1) - Q_k(n) = A_k(n) - x_k(n+1)$

The following inequality holds in both cases (by inspection):

$$Q_k(n)(Q_k(n+1) - Q_k(n)) \leq Q_k(n)(A_k(n) - x_k(n+1)).$$

Also, $\|Q(n+1) - Q(n)\|^2 \leq K$ since $|A_k(n) - x_k(n+1)| \leq 1$ for all k , so $\|Q(n+1) - Q(n)\|^2 \leq K$. Hence:

$$\|Q(n+1)\|^2 \leq \|Q(n)\|^2 + 2\langle Q(n), A(n) - x(n+1) \rangle + K.$$

Talking conditional expectations, since $x(n+1)$ is a function of $Q(n)$:

$$\mathbb{E}[\|Q(n+1)\|^2 - \|Q(n)\|^2 | Q(n)] \leq 2\langle Q(n), \lambda - x(n+1) \rangle + K.$$

Define $e = (1, \dots, 1)$. Since λ is in the interior of Λ , there exists $\epsilon > 0$ such that $\lambda + e\epsilon \in \Lambda$. Hence there exists $\alpha \in A$ such that $\lambda + e\epsilon = \sum_{i=1}^I \alpha_i x^i$. By definition of $x(n+1)$, for all x^i :

$$\langle Q(n), x(n+1) \rangle = \max_{x \in \mathcal{I}} \langle Q(n), x \rangle \geq \langle Q(n), x^i \rangle.$$

so:

$$\begin{aligned} \langle Q(n), x(n+1) \rangle &\geq \sum_{i=1}^I \alpha_i \langle Q(n), x^i \rangle \\ &= \langle Q(n), \lambda + e\epsilon \rangle \\ &= \langle Q(n), \lambda \rangle + \epsilon \sum_{k=1}^K Q_k(n). \end{aligned}$$

Going back:

$$\mathbb{E}[\|Q(n+1)\|^2 - \|Q(n)\|^2 | Q(n)] \leq -2\epsilon \sum_{k=1}^K Q_k(n) + K.$$

Define $V(\cdot) = \|\cdot\|^2$, and $F = \{q \in \mathbb{N}^K : 2\epsilon \sum_{k=1}^K q_k - K \geq \epsilon\}$. F is finite, and V is a Liapunov function since for all $q \notin F$:

$$\mathbb{E}[V(Q(n+1)) - V(Q(n)) | Q(n) = q] \leq -\epsilon.$$

By applying Foster's criterion, we have proven that $(Q(n))_n$ is positive recurrent which concludes the proof. \square

In fact a natural generalization of Max Weight can be obtained using the Liapunov function $V(Q(n)^\alpha)$, $\alpha > 1$, which yields the following algorithm:

$$x(n+1) = \arg \max_{x \in \mathcal{I}} \langle x, Q(n)^{\alpha-1} \rangle.$$

7.2.3 Computational complexity and message passing

We have established that Max-Weight is throughput optimal for any interference graph, which makes it a good candidate. However, Max-Weight suffers, by design, two important problems. At each time n , to calculate $x(n+1)$ we must find the maximal weighted independent set of weighted graph $G = (V, E, Q(n))$. This poses two problems when K is large:

- Computational complexity: finding the maximum weighed independent set of an arbitrary graph provably requires an exponential number of operations in the number of links K . So Max-Weight only applies to small networks, or particular topologies, such as matchings.
- Message passing: Max-Weight is a *centralized* algorithm, i.e at time n , to decide whether on not to activate link k , knowing $Q_k(n)$ is not sufficient, and one must know the state of all links $Q(n) = (Q_1(n), \dots, Q_K(n))$.

7.3 Distributed scheduling: CSMA with Glauber Dynamics

7.3.1 The algorithm

We now study an algorithm which is throughput optimal and fully distributed. Namely there is no message passing, and $x_k(n+1)$ is a function of $Q_k(n)$. This solves the message passing issue. The algorithm is CSMA-like, namely at each time n each transmitter decides to access the medium with a given probability, which might depend on $Q_k(n)$, but is independent of $Q_{k'}(n)$ for all $k' \neq k$. We define e_k the k -th canonical basis vector. Consider the following procedure.

- At time n , a user k is chosen uniformly at random
- (free the channel) If $x_k(n) = 1$ then:
 - $x(n) = x(n) - e_k$ with probability a_k
 - $x(n+1) = x(n)$ with probability $1 - a_k$.
- (seize the channel) If $x_k(n) = 0$, and $x(n) + e_k \in \mathcal{I}$ then
 - $x(n+1) = x(n) + e_k$ with probability b_k
 - $x(n+1) = x(n)$ with probability $1 - b_k$.
- (non-admissible schedule) If $x_k(n) = 0$ and $x(n) + e_k \notin \mathcal{I}$ then $x(n+1) = x(n)$.

Note that picking a user uniformly at random might be done using Poisson clocks. We define $r_k = \log(a_k/b_k)$ which is the “transmission aggressiveness”, and $r = (r_k)_k$ the corresponding vector.

Proposition 21 For each $r \in \mathbb{R}^K$, $(x(n))_n$ is a reversible Markov chain on \mathcal{I} with stationary distribution $(\pi_x)_{x \in \mathcal{I}}$:

$$\pi_x(r) = \mathbb{P}[x(n) = x] = \frac{\exp(\langle x, r \rangle)}{\sum_{x' \in \mathcal{I}} \exp(\langle x', r \rangle)}$$

Proof. It is sufficient to check that detailed balance holds. Define the transition probabilities:

$$P_{x,x'} = \mathbb{P}[x(n+1) = x' | x(n) = x].$$

It should be noted that $P_{x,x'}$ iff there exists k such that $x_{k'} = x'_{k'}$ for all $k' \neq k$. Consider k and $x \in \mathcal{I}$ such that $x_k = 0$ and $x + e_k \in \mathcal{I}$. Then:

$$\begin{aligned} P_{x,x+e_k} &= a_k/K \\ P_{x+e_k,x} &= b_k/K. \end{aligned}$$

Hence:

$$\frac{P_{x,x+e_k}}{P_{x+e_k,x}} = a_k/b_k = \exp(r_k).$$

By inspection of the formula given for π

$$\frac{\pi_{x+e_k}(r)}{\pi_x(r)} = \exp(r_k) = \frac{P_{x,x+e_k}}{P_{x+e_k,x}}.$$

Hence detailed balance holds, concluding the proof. \square

7.3.2 Throughput Optimality

First consider both transmission probabilities r and arrival rates λ fixed. The probability that link k is active is given by:

$$s_k(r) = \mathbb{E}[x_k(n)] = \sum_{x \in \mathcal{I}} x_k \pi_x(r).$$

Hence the network is stable for parameter r iff $s_k(r) > \lambda_k$ for all k . Consider λ in the interior of Λ . Then there exists $\alpha \in A$ such that $\lambda = \sum_{i=1}^I \alpha_i x^i$ and $\alpha_i > 0$. If we could set r such that $\pi_x(r) = \alpha_i$ for all i then we would obtain a stable configuration, in the sense that $s_k(r) \geq \lambda_k$ for all k . It is hence natural to minimize the distance between vectors $\pi(r)$ and α . Since they are both probability distributions the most natural metric is the Kullback-Leibler divergence:

$$F(r) = \sum_{i=1}^I \alpha_i \log(\alpha_i / \pi_{x^i}(r)).$$

We now prove that $F(r)$ has a unique, finite maximizer r^* and that setting $r = r^*$ guarantees stability of the network.

Proposition 22 Consider λ in the interior of Λ . Then:

- (i) Function $r \rightarrow F(r)$ is convex on $(\mathbb{R}^+)^K$.
- (ii) F has a unique, finite maximizer in $(\mathbb{R}^+)^K$ denoted r^*
- (iii) For all k : $s_k(r) \geq \lambda_k$, i.e. under transmission parameters r , the network is stable for any arrival rates strictly smaller than λ .

Proof.

- (i) Let us rewrite F as follows:

$$\begin{aligned} F(r) &= \sum_{i=1}^I \alpha_i \log(\alpha_i / \pi_{x^i}(r)) \\ &= \sum_{i=1}^I \alpha_i \log(\alpha_i) - \sum_{i=1}^I \alpha_i \log(\pi_{x^i}(r)). \end{aligned}$$

Furthermore, replacing π by its expression:

$$\begin{aligned} \sum_{i=1}^I \alpha_i \log(\pi_{x^i}(r)) &= \sum_{i=1}^I \left[\alpha_i \langle x^i, r \rangle - \alpha_i \log \left(\sum_{i'=1}^I \exp(\langle x^{i'}, r \rangle) \right) \right] \\ &= \langle \lambda, r \rangle - \log \left(\sum_{i=1}^I \exp(\langle x^i, r \rangle) \right). \end{aligned}$$

So:

$$F(r) = \sum_{i=1}^I \alpha_i \log(\alpha_i) - \langle \lambda, r \rangle + \log \left(\sum_{i=1}^I \exp(\langle x^i, r \rangle) \right).$$

Then $r \mapsto \langle \lambda, r \rangle$ is linear and $r \mapsto \log(\sum_{i=1}^I \exp(\langle x^i, r \rangle))$ is a log-sum-exp function, hence convex (see for instance [6][page 88]). We have proven that $r \mapsto F(r)$ is convex.

- (ii) Define $F^* = \inf_{r \geq 0} F(r)$. Since F is convex and positive, we either have that:

- (a) if a maximizer r^* exists then it is unique
- (b) otherwise there exists a sequence $(r^n)_n$ such that $F(r^n) \xrightarrow{n \rightarrow \infty} F^*$ and $\|r^n\| \xrightarrow{n \rightarrow \infty} \infty$.

We proceed by contradiction. Consider that case (b) occurs. Then $r^n / \|r^n\|$ is bounded, and denote by \bar{r} one of its accumulation points. Consider $y \geq 0$, then $F^* = \lim_{y \rightarrow \infty} F(y\bar{r})$. Define $m = \max_i \langle x^i, \bar{r} \rangle$ and denote by $J = \{i : \langle x^i, \bar{r} \rangle = m\}$. then $\pi(y\bar{r})$ tends to the uniform distribution on J . Indeed:

$$\pi_{x^i}(y\bar{r}) = \frac{\exp(y \langle x^i, \bar{r} \rangle)}{\sum_{i'=1}^I \exp(y \langle x^{i'}, \bar{r} \rangle)} = \frac{\exp(y(\langle x^i, \bar{r} \rangle - m))}{\sum_{i'=1}^I \exp(y(\langle x^{i'}, \bar{r} \rangle - m))} \xrightarrow{y \rightarrow \infty} \frac{\mathbf{1}\{i \in J\}}{|J|}$$

Now distinguish two cases:

$$J = \{1, \dots, N\}$$

In that case, for all y we have that $i \mapsto \langle x^i, \bar{r} \rangle$ is a constant, hence $\pi_{x^i}(y\bar{r}) = \pi_{x^i}(0)$ and $F(y\bar{r}) = F(0)$. Hence $F^* = F(0)$ so that F has a finite maximizer, a contradiction.

$$J \neq \{1, \dots, N\}$$

In that case there exists $i \notin J$ so that $\pi_{x^i}(y\bar{r}) \xrightarrow[n \rightarrow \infty]{} 0$. In turn $F^* = \lim_{y \rightarrow \infty} F(y\bar{r}) = +\infty$ since we have assumed that $\alpha_i > 0$. This is also a contradiction since $F(0) < +\infty$.

In summary, we have proven that F has a maximizer r^* , and this maximizer must be unique since F is a convex function.

(iii) Now consider r^* the unique maximizer of F in $(\mathbb{R}^+)^K$. Differentiating we obtain:

$$\frac{dF(r)}{dr_k} = s_k(r) - \lambda_k.$$

r^* maximizes F on $(\mathbb{R}^+)^K$ the constraints $r \geq 0$, so by the Karush-Kuhn-Tucker conditions this implies that there exists a positive vector d such that $\lambda - s(r^*) + d = 0$, hence $s(r^*) \geq \lambda$, which is the announced result. \square

7.3.3 Iterative scheme

We may now tune r sequentially, denote by $r(n)$ the value of the transmission parameters at time n . Assume that, when r is set equal to $r(n)$, link k is able to observe unbiased estimates of the throughput $s_k(r(n)) + M_k(n)$ and of the arrival rate $\lambda_k + M'_k(n)$ where $(M(n))_n$ and $(M'(n))_n$ are random sequences with null expectation. Such estimates can be obtained by calculating the empirical throughput and arrival rate over a suitably large time window. Consider the iterative scheme:

$$r_k(n+1) = r_k(n) + \epsilon_n [\lambda_k - s_k(r(n)) + M'_k(n) - M_k(n)].$$

with $(\epsilon_n)_n$ a sequence of positive step sizes. The above scheme is exactly a stochastic approximation scheme, and is ensured to converge under standard assumptions for stochastic approximation. Indeed the associated ODE is precisely:

$$\dot{r} = \lambda_k - s(r) = -\nabla F(r).$$

and admits the Liapunov function F . Furthermore, the above scheme is *fully distributed*, namely link k adjusts her transmission parameter $r_k(n)$ by observing her own throughput and arrival rate. There is no message passing between the links.

7.4 References

The Max-Weight policy (and a generalization known as Backpressure) are due to [28]. Max-Weight and its variants have been the subject of extensive research, and a comprehensive overview is given in [12]. An important contribution is [27], proving that Max-Weight is optimal (minimizes the expected packet delay) at the stability limit [27]. In fact, [27] proves that a phenomenon known as *state space collapse* occurs, namely, with high probability, the congestion process $Q(n)$ lies in a one dimensional subspace, irrespective of the number of links. The distributed CSMA algorithm described above was proposed in [16]

Chapter 8

Capacity scaling of wireless networks

Unlike the previous section, where we were essentially concerned by the performance of wireless networks at the PHY and MAC level, we now adopt a macroscopic view, where we consider a large wireless network with n nodes, and we are concerned with the network throughput (the number of successfully transmitted bits per unit time) when $n \rightarrow \infty$. The rationale is to determine whether or not the network throughput scales linearly in n . If it is the case, the network is scalable, in the sense that the throughput of a node is at least greater than a constant when $n \rightarrow \infty$. Otherwise the network is not scalable, so that the throughput of any individual node vanishes when $n \rightarrow \infty$, and it is not reasonable to operate the network in that regime. We follow the analysis of the seminal paper [14], and the reader may refer to the large body of literature building on [14]. One of the most notable follow-ups to [14] is [13] where node mobility is proven to have a critical impact on capacity scaling.

8.1 The model

8.1.1 Node locations

Let us describe the model. We consider \mathcal{X} a set of unit surface. We will either consider \mathcal{X} to be a disk of unit area on the plane, or \mathcal{X} to be a sphere of unit surface in 3-dimensional space. In both situations we will denote by $|x - x'|$ the distance between two points of \mathcal{X} . If \mathcal{X} is a disk on the plane, $|\cdot|$ denotes the Euclidean distance, and if \mathcal{X} is a sphere in 3-dimensional space, then $|\cdot|$ denotes the length of the shortest path on the sphere between x and x' . We consider n nodes with positions $(x_1, \dots, x_n) \in \mathcal{X}$. We assume that a bandwidth of unit size is shared between the n nodes. This is without loss of generality since, in this model, the capacity lower and upper bounds are linear functions of the available bandwidth.

We consider a mapping which associates each node with the corresponding destina-

tion $d : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$, where $d(i) \neq i$ denotes the destination of node i . In all cases we will assume that interference is described by the protocol model. Namely node i has a transmission range r_i , and she may communicate with its destination $d(i)$ iff they are at a distance smaller than r_i , and no other transmitting node is located at a distance less than r'_i from $d(i)$ with $r'_i > r_i$.

8.1.2 Transmission schedules

Let us now describe what are the possible subsets of nodes that may transmit simultaneously. We first describe the case where nodes apply power control, so that their transmission range coincides with the distance at which their intended receiver is placed. We define a schedule: $s = (X_1, Y_1), \dots, (X_I, Y_I)$, with I transmitters and where X_i, Y_i denotes the i -th transmission/receiver pair. Schedule s is feasible iff for all i and $j \neq i$:

$$|X_j - X_i| > (1 + \Delta)|X_i - Y_i|$$

Therefore, in this model, exclusion regions are circles. An exclusion region is a region centred on a given node in which no other node may transmit simultaneously. When a feasible $s = (X_1, Y_1), \dots, (X_I, Y_I)$ is selected during a time slot, receivers X_i sends one packet to receiver Y_i , for all i . Now let us denote by \mathcal{T} set of feasible schedules. A transmission strategy is a sequence of feasible schedules. Consider the situation where packets arrive at rate λ at each node. We say that arrival rate λ is feasible iff, there exists a transmission strategy such that the corresponding queuing system is stable. λ is feasible iff there exists a strategy such that for all i , nodes i and $d(i)$ can exchange λ packets per unit of time.

We define λ^* the largest feasible λ , and we define the capacity $C(n) = n\lambda^*$. The capacity is the maximal amount of information that can be exchanged through the network, with a fairness constraint that all nodes get equal throughput. A related measure of performance is the transport capacity $T(n)$ which is defined as the number of bits/meters that can be transported by the network in a unit of time. It is calculated as the sum of the travelled distance by all the transmitted packets in a unit of time.

8.1.3 High probability events

We will consider two possible setting for the node positions and the destination mapping:

- Arbitrary networks: positions $x^n = (x_1, \dots, x_n) \in \mathcal{X}$ and the destination mapping d are fixed.
- Random networks: positions $x^n = (x_1, \dots, x_n) \in \mathcal{X}$ are i.i.d uniformly distributed on \mathcal{X} and the destination mapping d is uniformly distributed on the set of permutations.

In the case of random networks, some configurations only occur when n is small, and appear with vanishing probability when $n \rightarrow \infty$. For instance, when n is large, the nodes must be evenly distributed on \mathcal{X} . We will use the following “high probability reasoning”. Namely, consider S_n, D_n a set of positions and destination mappings. We will say that (S_n, D_n) occurs with high probability iff $\mathbb{P}[x^n \in S_n, d^n \in D_n] \xrightarrow{n \rightarrow \infty} 1$. High probability reasoning allows us to consider only “typical networks” rather than arbitrary ones.

8.2 Arbitrary networks on a circle

We start our study by considering arbitrary networks, and \mathcal{X} denotes a sphere with unit surface in the plane.

8.2.1 Capacity upper bound

The first result gives an upper bound on the capacity achievable by any transmission policy. The intuitive explanation for this upper bound is that transmissions “consume space”, so that any receiver generates an exclusion zone around her, in which no other transmission can take place. Hence the maximum amount of transmissions that can take place is closely related to sphere packing. The proof technique involves upper bounding the transport capacity.

Proposition 23 Define $L = (1/n) \sum_{n'=1}^N |x_i - x_{d(i)}|$ the mean length of a connection. The capacity admits the upper bound:

$$C(n) \leq \frac{\sqrt{8n}}{\sqrt{\pi}L}.$$

Proof. The proof technique involves upper bounding the transport capacity rather than the capacity. Consider a schedule $s = (X_1, Y_1), \dots, (X_I, Y_I)$ used during a time slot. Define $r_i = |X_i - Y_i|$ the receiver-transmitter distance for the i -th active link. Denote by T the amount of transported information when schedule s is used during a time slot. We upper bound T , using the fact that there are at most $n/2$ active transmitters and the Cauchy-Schwartz inequality:

$$T = \sum_{i=1}^I r_i \leq \sqrt{I} \left(\sum_{i=1}^I r_i^2 \right)^{1/2} \leq \sqrt{n/2} \left(\sum_{i=1}^I r_i^2 \right)^{1/2}$$

Schedule s must be feasible, so that the following constraints are satisfied. For all i, j with $i \neq j$ we have:

$$(1 + \Delta)|X_i - Y_i| \leq |X_j - Y_i| \leq |X_j - Y_j| + |Y_j - Y_i|$$

so that:

$$(1 + \Delta)|X_i - Y_i| - |X_j - Y_j| \leq |Y_j - Y_i|,$$

using the triangle inequality. Exchanging the roles of i and j in the above reasoning:

$$(1 + \Delta)|X_j - Y_j| - |X_i - Y_i| \leq |Y_i - Y_j|,$$

and adding the inequalities we get:

$$(\Delta/2)(r_i + r_j) \leq |Y_i - Y_j|.$$

The above inequality has the following geometric interpretation. For all i , define \mathcal{C}_i a disk of radius $r_i\Delta/2$ centred at Y_i . Then for all $i \neq j$ we have $\mathcal{C}_i \cap \mathcal{C}_j = \emptyset$. Denote by $A(\cdot)$ the area of a set. We have:

$$1 = A(\mathcal{S}) \geq \sum_{i=1}^I A(\mathcal{C}_i \cap \mathcal{S}).$$

Since $Y_i \in \mathcal{S}$, we have $A(\mathcal{C}_i \cap \mathcal{S}) \geq A(\mathcal{C}_i)/4 = \pi\Delta^2 r_i^2/16$. Hence:

$$\left(\sum_{i=1}^I r_i^2\right)^{1/2} \leq 4/(\sqrt{\pi}\Delta).$$

Replacing:

$$T \leq \frac{\sqrt{n}}{\sqrt{2}} \frac{4}{\sqrt{\pi}\Delta} = \frac{\sqrt{8n}}{\sqrt{\pi}\Delta}.$$

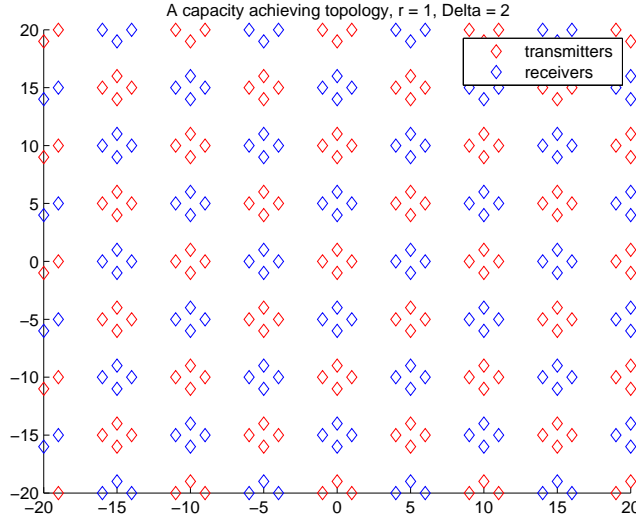
The above reasoning is valid for any feasible schedule. Hence:

$$C(n)L = T(n) \leq \frac{\sqrt{8n}}{\sqrt{\pi}\Delta}.$$

which concludes the proof. □

8.2.2 Capacity lower bound

Let us now exhibit a topology where the capacity of the network scales as the upper bound. Since we are mainly concerned with the regime $n \rightarrow \infty$ we may consider n a multiple of 4. The topology of choice is represented by the following picture:



Namely, for all $(j, k) \in \mathbb{Z}^2$, define $a(j, k) = r(1 + 2\Delta)(j, k)$. Then, at locations $a(i, j) + (\Delta r, 0)$, $a(i, j) - (\Delta r, 0)$, $a(i, j) + (0, \Delta r)$, $a(i, j) - (0, \Delta r)$, place transmitters if $|j + k|$ is even, and receivers if $|j + k|$ is odd.

Proposition 24 Consider the above network with transmission range chosen as:

$$r^{-1} = (1 + 2\Delta)(\sqrt{n/4} + \sqrt{2\pi})$$

The transmission scheme where all transmitters are always active and transmit to their nearest neighbour achieves a capacity of:

$$C(n) = \frac{n}{2L(1 + 2\Delta)(\sqrt{n/4} + \sqrt{2\pi})}.$$

Proof. We only provide a sketch of proof. Let us check that the proposed transmission scheme is admissible/feasible. Consider a given receiver. The distance to the nearest transmitter is $d = r$. The distance to the second nearest transmitter is:

$$\bar{d} = \sqrt{(\Delta r)^2 + r^2(1 + \Delta)^2} \geq r(1 + \Delta) = d(1 + \Delta).$$

Hence the schedule is feasible. We may check that one can pack \sqrt{nr} transmitter-receiver pairs in the region of interest \mathcal{X} . There are $n/2$ simultaneous transmissions of range r , and hence the transport capacity is:

$$T(n) = nr/2,$$

which is the announced result, replacing r by its expression. □

8.3 Random networks on a sphere

We turn to random networks on the sphere in 3-dimensional space. Here \mathcal{X} is the sphere with unit surface in \mathbb{R}^3 . We do not consider power control, so that all nodes transmit with fixed power and have a fixed range r . Therefore, for any feasible scheme $s = (X_1, Y_1), \dots, (X_I, Y_I)$, we must have: $|X_i - Y_i| \leq r$ for all i and $|Y_j - X_i| \geq (1 + \Delta)r$ for all $i \neq j$. We denote by $A(r)$ the area of a disk of radius r on the sphere.

8.3.1 Connectivity

Since the network is random, and nodes are not placed in a regular manner, the first problem is to select the transmission range $r(n)$ such that there are no isolated nodes with high probability. Isolated nodes are nodes that are not within the transmission range of any other node. Of course, if there exists an isolated node, then she may not send or receive any data, and the achievable network throughput is null. Let us define the graph $\mathcal{G}(n) = (V, E)$ with $V = \{1, \dots, n\}$ and (x_i, x_j) iff $|x_i - x_j| \leq r(n)$. We give a necessary condition and a sufficient condition for connectivity. Essentially, one may show that $r(n)^2$ should be of order $\log(n)/n$ to ensure connectivity with high probability. Therefore, we assume that $r(n)$ is of order $\sqrt{\frac{\log n}{n}}$.

Proposition 25 *If $r(n) \leq \sqrt{\frac{k}{n\pi}}$ with $k > 0$ a constant, then \mathcal{G}_n is not connected with high probability.*

Proof. Define P the probability that there exists an isolated node, and P_i probability that i is isolated.

(i) Nodes are i.i.d uniformly distributed on the sphere with unit surface. Therefore the probability that a given node falls in a set is equal to its surface. P_i is the probability that no other node falls within a disk of radius r centred at the location of i . Hence we have $P_i = (1 - A(r(n)))^{n-1}$. Now:

$$P > P_1 = (1 - A(r(n)))^{n-1} \geq (1 - \pi r(n)^2)^{n-1} \sim (1 - k/n)^{n-1} \xrightarrow[n \rightarrow \infty]{} e^{-k}.$$

Hence P does not vanish when $n \rightarrow \infty$, and \mathcal{G}_n cannot be connected with high probability. □

8.3.2 Source destination paths

Recall that, in the random network model, source-destination pairs are chosen randomly independent of node locations. The path between a transmitter and a receiver is the great circle between their respective location, and for source-destination pair i , we denote the corresponding path by $L_i \subset \mathcal{X}$.

Proposition 26 $(L_1, \dots, L_{n/2})$ are i.i.d. Define $l = \mathbb{E}[L_1]$ and $\bar{L} = (2/n) \sum_{i=1}^{n/2} L_i$. Then $\bar{L} \geq l/2 > 0$ with high probability.

8.3.3 Upper bound on the capacity

We may now state an upper bound on the achievable capacity for random networks on a sphere. The arguments are very similar to those used for arbitrary networks in the plane: active transmissions consume space, so that sphere packing gives an upper bound on the amount of transported information for any feasible schedule. The upper bound decreases as a function of $r(n)$ so that $r(n)$ should be chosen as small as possible. On the other hand $r(n)$ must be larger than $\mathcal{O}(\frac{\log n}{n})$ to ensure connectivity, so that $C(n)$ can be (at most) $\mathcal{O}(\sqrt{\frac{n}{\log n}})$.

Proposition 27 The capacity is upper bounded by $C(n) \leq \bar{C}(n)$ with high probability, with

$$\bar{C}(n) \sim \frac{8}{l\pi\Delta^2 r(n)}, \quad n \rightarrow \infty$$

Proof. Consider a feasible schedule with N transmitters-receiver pairs. Define \mathcal{C}_i the disk of radius $\Delta r(n)/2$ centred at the i -th transmitter. By the same reasoning as above, the disks $(\mathcal{C}_i)_i$ are disjoint, and are a subset of the sphere with unit surface. Hence $NA(r(n)) \leq 1$ and

$$N \leq \frac{1}{A(\Delta r(n)/2)}.$$

The distance travelled by the N transmitted packets $T(n)$ verifies:

$$T(n) \leq Nr(n) \leq \frac{r(n)}{A(\Delta r(n)/2)} \sim \frac{4}{\pi\Delta^2 r(n)},$$

using the fact that $A(r) \sim \pi r^2$ when $r \rightarrow 0$.

The capacity verifies, with high probability:

$$C(n) = \frac{T(n)}{L} \leq \frac{2T(n)}{l}.$$

which is the announced result. □

8.3.4 Lower bound on the capacity

We now propose a scheme which achieves a throughput of $\mathcal{O}(\sqrt{\frac{n}{\log n}})$, matching the upper bound on the capacity (up to a multiplicative factor). The proof is long, and is split in several subsections. Essentially, each part of the proof involves showing that, with high probability, a certain type of configuration arises.

8.3.5 Partitioning the domain into cells

First we consider a partition of \mathcal{X} : $\mathcal{V}_n = (V_1, \dots, V_N)$ with the following properties. Define $a(n) = 100 \log(n)/n$ and $\rho(n)$ such that $A(\rho(n)) = a(n)$. We assume that for all i and $j \neq i$ the following holds:

- V_i contains a disk of radius $\rho(n)$
- V_i is contained in a disk of radius $2\rho(n)$

Proposition 28 *There exists such a partition.*

Proof. Define a sequence of points a_1, \dots, a_N the following way. Choose a_1 arbitrarily in the sphere. For all $i \geq 0$, if there exists $x \in \mathcal{S}$ such that $D(x, \rho(n))$ does not intersect with $D(a_i, \rho(n))$ for all $1 \leq i < n$ then set $a_i = x$. Otherwise $n = N$ and the procedure stops. The procedure must finish in finite time since $1 \geq N100 \log(n)/n$. Then return (V_1, \dots, V_N) the Voronoi cells of (a_1, \dots, a_N) . Since $|a_i - a_j| \geq 2\rho(n)$ for all $i \neq j$, V_i contains a disk of radius $\rho(n)$. Consider x and a_i the closest point to x . Assume that $|x - a_i| \geq 2\rho(n)$. Then $|x - a_j| \geq 2\rho(n)$ for all j and x so that $\{a_1, \dots, a_N\}$ is not maximal. \square

We may now define adjacency and interference for cells close to each other:

- V_i and V_j interfere if there exists $x_i \in V_i, x_j \in V_j$ such that

$$|x_i - x_j| \leq r(n)(2 + \Delta).$$

- V_i and V_j are adjacent if $V_i \cap V_j \neq \emptyset$.

We choose $r(n) = 8\rho(n)$ which ensures that all adjacent cells communicate.

8.3.6 Schedules

We now turn to schedules. The idea is to choose a subset of cells to activate, and, for each activated cell select a transmitter to activate.

Proposition 29 (i) *Each cell has at most $I = \mathcal{O}(\Delta^2)$ interferers.*

(ii) *There exists a schedule of length $I + 1$ such that each cell transmits at least a packet.*

Proof. (i) Consider V_j interfering with V_i . Then there exists $x_i \in V_i, x_j \in V_j$ such that $|x_i - x_j| \leq r(n)(2 + \Delta)$. Consider $y_j \in V_j$. Each cell has diameter of at most $4\rho(n)$. Hence:

$$|a_i - y_j| \leq |a_i - x_i| + |x_i - x_j| + |x_j - y_j| \leq 6\rho(n) + (2 + \Delta)r(n)$$

Hence all interferers of V_i are comprised in a disk of radius $6\rho(n) + (2 + \Delta)r(n)$ centered at a_i . The area of a cell is at least $\pi\rho(n)^2$, hence N_i , the number of interferers of V_i verifies:

$$\pi\rho(n)^2 N_i \leq \pi(6\rho(n) + (2 + \Delta)r(n))^2$$

Hence:

$$N_i \leq \frac{(6\rho(n) + (2 + \Delta)r(n))^2}{\rho(n)^2} = 2(11 + 4\Delta^2).$$

since $r(n) = 8\rho(n)$. Hence $N_i = \mathcal{O}(\Delta^2)$.

(ii) Define the graph G whose vertices are the cells, and where there is an edge between two cells if and only they interfere. The maximal degree of G is I , hence it can be coloured with $I + 1$ colors, such that the colors of any two neighbouring vertices are distinct. Then consider the schedule of length $I + 1$ where at times $i \in \{1, \dots, I + 1\}$, only vertices of color number i are activated. This schedule lets all cells transmit one packet, never lets two interfering cells transmit simultaneously. \square

8.3.7 Each cell has at least one node whp

We now establish that each cell contains at least one node with high probability.

Proposition 30 *With high probability, V_i contains a node for all i .*

Proof. Define P_i the probability that V_i is empty and P the probability that there exists i such that V_i is empty. V_i contains at least a disk of area $a(n)$ so:

$$P_i \leq (1 - a(n))^{n-1}.$$

Each cell has area at least $a(n)$ so that there are less than $N \leq 1/a(n) \leq n$ cells. Using a union bound:

$$P \leq \sum_{i=1}^N P_i \leq n(1 - a(n))^{n-1}$$

Taking logarithms:

$$\log(P) = \log(n) + (n-1) \log(1 - a(n)) = \log(n) - 100 \log(n) \frac{n-1}{n} + o(1) \xrightarrow{n \rightarrow \infty} -\infty.$$

Hence $P \xrightarrow{n \rightarrow \infty} 0$ so that each cell has at least one node with high probability. \square

8.3.8 Routing

Recall that L_i is a straight line (a great circle) linking the i -th transmitter pair. We now propose the following routing scheme. Define $V^i(1), V^i(2), \dots, V^i(R_i)$ the successive cells traversed by L_i . Then packets are transferred through R_i hops: $X_i \rightarrow V^i(1) \rightarrow V^i(2) \rightarrow \dots \rightarrow V^i(R_i) \rightarrow Y_i$. The only remaining element is to upper bound the number of routes traversing a given cell.

Proposition 31 *With high probability, the maximal number of lines going through any cell is at most $\mathcal{O}(\sqrt{n \log(n)})$.*

Proof. The proof is beyond the scope of this lecture notes and relies on a concept called the Vapnik-Chervronakis (VC) dimension. Given a set of sets (also called a class), the VC dimension measures its complexity in terms of “shattering”. The VC dimension is ubiquitous in statistical learning, and essentially allows to derive a uniform version of the weak law of large numbers. The interested reader can refer to, for instance [32]. \square

8.3.9 Capacity

Proposition 32 *With high probability, the scheme described above has a capacity of at least:*

$$C(n) \geq \mathcal{O} \left((1 + \Delta)^{-2} \sqrt{\frac{n}{\log n}} \right).$$

Proof. For any transmitter-receiver pair, any packet goes through any cell at most once. So it is sufficient to ensure that the rate at which packets depart all cells is greater than the corresponding arrival rate. Denote by $\lambda(n) = \frac{C(n)}{n}$ the maximal admissible arrival rate. Consider cell V_i . With high probability, the expected number of packets entering V_i per unit of time is $\mathcal{O}(\lambda(n)\sqrt{n \log n})$. On the other hand, cell V_i transmits a packet (at least) every $I + 1 = \mathcal{O}((1 + \Delta)^2)$ time slots. Hence each cell is a server with arrival rate $\mathcal{O}(\lambda(n)\sqrt{n \log n})$, and service rate $\mathcal{O}((1 + \Delta)^2)$. Hence, with high probability, the capacity is at least $\mathcal{O} \left((1 + \Delta)^{-2} \sqrt{\frac{n}{\log n}} \right)$. \square

8.4 Area of a disk on a sphere

We use repeatedly the following inequalities relating the area of a disk on the sphere to that of a disk on the two dimensional plane.

Proposition 33 *We have:*

$$1 - \frac{\pi r^2}{3} \leq \frac{A(r)}{\pi r^2} \leq 1.$$

Bibliography

- [1] N. Abramson. The throughput of packet broadcasting channels. *IEEE Trans. on Communications*,, 1977.
- [2] F. Baccelli and P. Bremaud. *Elements of Queueing Theory*. Springer Verlag, 2003.
- [3] G. Bianchi. Performance analysis of the ieee 802.11 distributed coordination function. *IEEE JSAC*, 2000.
- [4] C. Bordenave, D. McDonald, and A. Proutiere. Asymptotic stability region of slotted aloha. *IEEE Trans. Information Theory*, 2012.
- [5] V. S. Borkar. *Stochastic approximation: a dynamical systems viewpoint*. Cambridge University Press, 2008.
- [6] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [7] P. Bremaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer, 2008.
- [8] G. Brown. Iterative solutions of games by fictitious play. *Activity Analysis of Production and Allocation*, 1951.
- [9] F. G. Foster. On the stochastic matrices associated with certain queuing processes. *Ann. Math. Statistics*, 1953.
- [10] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press Books. The MIT Press, July 1998.
- [11] M. S. Gast. *802.11 Wireless Networks - The Definitive Guide*. O'REILLY, 2008.
- [12] L. Georgiadis, M. J. Neely, and L. Tassiulas. *Resource Allocation and Cross-Layer Control in Wireless Networks*. Foundations and Trends in Networking, 2006.
- [13] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. on Networking*, 2002.

- [14] P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Trans. on Information Theory*, 2000.
- [15] IEEE. 802.11 standard. <http://standards.ieee.org/about/get/802/802.11.html>.
- [16] L. Jiang and J. Walrand. A distributed csma algorithm for throughput and utility maximization in wireless networks. *Proc. of Allerton*, 2008.
- [17] F. Kelly. *Reversibility and Stochastic Networks*. Cambridge University Press, 1979.
- [18] J. Kiefer and J. Wolfowitz. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):pp. 462–466, 1952.
- [19] A. Kumar, E. Altman, E. Miorandi, and M. Goyal. New insights from a fixed point analysis of single cell ieee 802.11 wlans. *IEEE INFOCOM*, 2005.
- [20] D. Levin, Y. Peres, and E. Vilmer. *Markov Chains and Mixing Times*. AMS, 2008.
- [21] A. Liapunov. The general problem of the stability of motion. (*In Russian*), 1892.
- [22] L. Ljung. Analysis of recursive stochastic algorithms. *Automatic Control, IEEE Transactions on*, 22(4):551–575, 1977.
- [23] S. Meyn and R. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, 2009.
- [24] E. Pardoux. *Processus de Markov et applications (in French)*. Dunod, 2006.
- [25] R. Rao and A. Ephremides. On the stability of interacting queues in a multiple-access system. *IEEE Trans. Information Theory*, 1988.
- [26] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, Sept. 1951.
- [27] A. Stolyar. Max-weight scheduling in a generalized switch: State-space collapse and workload minimization in heavy traffic. *Annals of Applied Probability*, 2004.
- [28] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks,. *IEEE Trans. on Automatic Control*, 1992.
- [29] G. Teschl. *Ordinary Differential Equations and Dynamical Systems*. AMS, 2012.

- [30] D. Tse and H. Vishwanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [31] B. Tsybakov and V. Mikhailov. Ergodicity of a slotted aloha system. *Problems Inform. Transmission*, 1979.
- [32] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 1971.
- [33] D. Williams. *Probability with Martingales*. Cambridge University Press, 1991.