



Thèse présentée pour obtenir le titre de

DOCTEUR DE L'ÉCOLE POLYTECHNIQUE

Spécialité : Mathématiques appliquées

par

Olivier Fercoq

**Optimization of Perron eigenvectors and applications:
From web ranking to chronotherapeutics**

**Optimisation de vecteurs propres de Perron et applications :
Du référencement de pages web à la chronothérapie**

Soutenue le 17 septembre 2012 devant le jury composé de :

Marianne Akian	INRIA Saclay et CMAP Ecole Polytechnique	co-directrice
Konstantin Avratchenkov	INRIA Sophia Antipolis	rapporteur
Mustapha Bouhtou	France Télécom R & D	co-directeur
Jean Clairambault	INRIA Rocquencourt et Université Paris 6	président du jury
Michel De Lara	Ecole des Ponts ParisTech et Université Paris-Est	examinateur
Stéphane Gaubert	INRIA Saclay et CMAP Ecole Polytechnique	directeur
Roberto Tempo	Politecnico di Torino	examinateur
Paul Van Dooren	Université Catholique de Louvain	rapporteur

Summary

Search engines play a key role in the World Wide Web. They gather information on the web pages and for each query of a web surfer they give a sorted list of relevant web pages. Internet search engines use a variety of algorithms to sort web pages based on their text content or on the hyperlink structure of the web. Here, we focus on algorithms that use the latter hyperlink structure, called link-based algorithms, among them PageRank, HITS, SALSA and HOTS. The basic notion for all these algorithms is the web graph, which is a digraph with a node for each web page and an arc between nodes i and j if there is a hyperlink from page i to page j .

The original problem considered in the present work, carried out as part of a collaboration between INRIA and Orange Labs, is the optimization of the ranking of the pages of a given web site. It consists in finding an optimal outlink strategy maximizing a scalar function of a given ranking subject to design constraints. PageRank, HITS and SALSA are Perron vector rankings, which means that they correspond to the principal eigenvector of a (elementwise) nonnegative matrix. When optimizing the ranking, we thus optimize a scalar utility function of the Perron eigenvector over a set of nonnegative irreducible matrices. The matrix is constructed from the web graph, so controlling the hyperlinks corresponds to controlling the matrix itself.

We first study general PageRank optimization problems with a “total income” utility function and design constraints. This case is of particular interest since the value of the PageRank is an acknowledged economic issue. We reduced the PageRank optimization problem to Markov decision problems such that the action sets are implicitly defined as the vertices of polytopes that have a polynomial time separation oracle. We show that such Markov decision problems are solvable in polynomial time and we provide a scalable algorithm for the effective resolution of the PageRank optimization problem on large dataset.

Then, we study the general problem of optimizing a scalar utility function of the Perron eigenvector over a set of nonnegative irreducible matrices. This covers all Perron vector rankings, including HITS and SALSA. We show that the matrix of partial derivatives of the objective has a low rank and can be computed by an algorithm with the same convergence properties as the power algorithm used to compute the ranking and so the value of the objective. We give an optimization algorithm that couples power and gradient iterations and prove its convergence to a stationary point of the optimization problem. Considering HOTS

as a nonlinear Perron vector, we show that the HOTS algorithm converges with a linear rate of convergence, that the objective of the HOTS optimization problem has a low rank and that the coupled power and gradient algorithm applies.

Finally, we extend the domain of application of the Perron eigenvalue and eigenvector optimization methods to the optimization of chemotherapy under the McKendrick model of population dynamics. We consider here that the cells behave differently at an hour of the day or another. We want to take advantage of this feature to minimize the growth rate of cancer cell population while we maintain the growth rate of healthy cell population over a given toxicity threshold. The objective and the constraint can be written as the Floquet eigenvalues of age-structured PDE models with periodic coefficients, and they are approximated by Perron eigenvalues in the discretized problem. We search for locally optimal drug infusion strategies by a method of multipliers, where the unconstrained minimizations are performed using the coupled power and gradient algorithm that we have developed in the context of web ranking optimization.

Résumé

Les moteurs de recherche jouent un rôle essentiel sur le Web. Ils rassemblent des informations sur les pages web et pour chaque requête d'un internaute, ils donnent une liste ordonnée de pages pertinentes. Ils utilisent divers algorithmes pour classer les pages en fonction de leur contenu textuel ou de la structure d'hyperlien du Web. Ici, nous nous concentrons sur les algorithmes qui utilisent cette structure d'hyperliens, comme le PageRank, HITS, SALSA et HOTS. La notion fondamentale pour tous ces algorithmes est le graphe du web. C'est le graphe orienté qui a un nœud pour chaque page web et un arc entre les nœuds i et j si il y a un hyperlien entre les pages i et j .

Le problème original considéré dans cette thèse, réalisée dans le cadre d'une collaboration entre INRIA et Orange Labs, est l'optimisation du référencement des pages d'un site web donné. Il consiste à trouver une stratégie optimale de liens qui maximise une fonction scalaire d'un classement donné sous des contraintes de design. Le PageRank, HITS et SALSA classent les pages par un vecteur de Perron, c'est-à-dire qu'ils correspondent au vecteur propre principal d'une matrice à coefficients positifs. Quand on optimise le référencement, on optimise donc une fonction scalaire du vecteur propre de Perron sur un ensemble de matrices positives irréductibles. La matrice est construite à partir du graphe du web, donc commander les hyperliens revient à commander la matrice elle-même.

Nous étudions d'abord un problème général d'optimisation du PageRank avec une fonction d'utilité correspondant au revenu total du site et des contraintes de design. Ce cas est d'un intérêt particulier car pour de nombreux sites la valeur du PageRank est corrélée au chiffre d'affaires. Nous avons réduit le problème d'optimisation du PageRank à des problèmes de décision markoviens dont les ensembles d'action sont définis implicitement comme étant les points extrêmes de polytopes qui ont un oracle de séparation polynomial. Nous montrons que de tels problèmes de décision markoviens sont solubles en temps polynomial et nous donnons un algorithme qui passe à l'échelle pour la résolution effective du problème d'optimisation du PageRank sur de grandes bases de données.

Ensuite, nous étudions le problème général de l'optimisation d'une fonction scalaire du vecteur propre de Perron sur un ensemble de matrices positives irréductibles. Cela couvre tous les classements par vecteur de Perron, HITS et SALSA compris. Nous montrons que la matrice des dérivées partielles de la fonction objectif a un petit rang et peut être calculée

par un algorithme qui a les mêmes propriétés de convergence que la méthode de la puissance utilisée pour calculer le classement. Nous donnons un algorithme d'optimisation qui couple les itérations puissance et gradient et nous prouvons sa convergence vers un point stationnaire du problème d'optimisation. En considérant HOTS comme un vecteur de Perron non linéaire, nous montrons que l'algorithme HOTS converge géométriquement et nous résolvons l'optimisation locale de HOTS.

Finalement, nous étendons le domaine d'application des méthodes d'optimisation du vecteur propre et de la valeur propre de Perron à l'optimisation de la chimiothérapie, sous l'hypothèse que les cellules se comportent différemment suivant l'heure de la journée. Nous voulons profiter de cette caractéristique pour minimiser le taux de croissance des cellules cancéreuses tout en maintenant le taux de croissance des cellules saines au dessus d'un seuil de toxicité donné. L'objectif et la contrainte peuvent s'écrire comme les valeurs propres de Floquet de modèles d'EDP structurés en âge avec des coefficients périodiques, qui sont approchés par des valeurs propres de Perron dans le problème discrétisé. Nous cherchons des stratégies d'injection de médicament localement optimales par une méthode des multiplicateurs où les minimisations sans contrainte sont faites en utilisant l'algorithme couplant les itérations puissance et gradient développé dans le cadre de l'optimisation du référencement.

Contents

1	Introduction	13
1.1	Context of the thesis	13
1.1.1	Web ranking	13
1.1.2	Age-structured population dynamics	15
1.2	Motivation	16
1.3	Contribution	17
1.4	Organization	20
1.5	Contexte de la thèse	23
1.5.1	Référencement web	23
1.5.2	Dynamiques de populations structurées en âge	25
1.6	Motivation	26
1.7	Contribution	28
1.8	Organisation	31
2	Web ranking and (nonlinear) Perron Frobenius theory	33
2.1	Google's PageRank	33
2.2	PageRank optimization in the literature	34
2.3	Other web ranking algorithms	37
2.4	Nonlinear Perron-Frobenius theory	40
3	Markov Decision Processes with implicit action spaces	43
3.1	Finite Markov Decision Processes	44
3.2	Polynomial time solvability of well-described Markov decision problems	45
3.3	Convex Markov decision processes	47
4	Ergodic control and Polyhedral approaches to PageRank optimization	51
4.1	Introduction	51
4.2	PageRank optimization problems	54
4.2.1	Optimization of PageRank	54

4.2.2	Design constraints of the webmaster	56
4.3	Reduction of the PageRank Optimization Problem with local constraints to Ergodic Control	57
4.3.1	Reduction of the Continuous PageRank Optimization Problem to Ergodic Control	57
4.3.2	The polytope of uniform transition measures	58
4.4	Solving the PageRank Optimization Problem with local constraints	61
4.4.1	Polynomial time solvability of the PageRank Optimization Problem with local constraints	61
4.4.2	Optimizing the PageRank via Value iteration	62
4.4.3	A generalization of the algorithm of Csáji et al. for total income utility functions	65
4.5	General shape of an optimized web site	67
4.5.1	Sensitivity of optimal solutions to incomplete crawling	69
4.6	PageRank Optimization with coupling constraints	70
4.6.1	Reduction of the problem with coupling constraints to constrained Markov decision processes	70
4.6.2	Optimality condition	72
4.6.3	A Lagrangian relaxation scheme to handle coupling constraints between pages	73
4.6.4	A heuristic to solve the discrete problem with coupling constraints	73
4.7	Experimental results	74
4.7.1	Continuous problem with local constraints only	75
4.7.2	Discrete problem	75
4.7.3	Coupling linear constraints	75
5	PageRank optimization applied to spam detection	77
5.1	Introduction	77
5.2	The MaxRank algorithm	78
5.3	Spam detection and PageRank demotion	83
6	Convergence of Tomlin's HOTS algorithm	87
6.1	Introduction	87
6.2	The ideal HOTS algorithm	89
6.3	The effective HOTS algorithm	93
6.4	An exact coordinate descent for the truncated scaling problem	102
6.5	Comparison with alternative algorithms	104
6.6	A normalized HOTS algorithm for a better convergence rate	106
7	Optimization of the Perron eigenvalue and eigenvector	109
7.1	Introduction	109
7.2	Perron vector and Perron value optimization problems	112
7.3	A power-type algorithm for the evaluation of the derivative of a function of the principal eigenvector	117
7.4	Coupling gradient and power iterations	122
7.5	Application to HITS optimization	127
7.6	Optimization of the eigenvector of a nonlinear Perron-Frobenius operator	130

7.7	Optimization of the HOTS score of a website's pages	132
7.8	Numerical results	135
7.8.1	HITS authority optimization	135
7.8.2	HOTS optimization	136
7.8.3	Scalability of the algorithms	137
7.9	Conclusion	137
8	Optimization of cancer drug treatments through age-structured population dynamics	139
8.1	Introduction	139
8.2	Drugs used in cancer treatments and their targets	141
8.2.1	Fate of drugs in the organism: molecular pharmacokinetics-pharmacodynamics	141
8.2.2	Cytotoxics and cytostatics	142
8.2.3	Representation of drug targets	142
8.3	Age-structured linear models for tissue proliferation and its control	143
8.3.1	Linear and nonlinear models	143
8.3.2	Age-structured models for tissue proliferation and its control	144
8.3.3	Basic facts about age-structured linear models	144
8.3.4	Discretization scheme	146
8.3.5	Numerical determination of the growth exponent	147
8.4	The control and its missions: representing the action of drugs	148
8.4.1	The control functions	148
8.4.2	Objective functions: measuring the output	148
8.4.3	Constraints, technological and biological, static or dynamic	149
8.4.4	An eigenvalue optimization problem with constraints	151
8.5	Identification of model parameters: the target model and drug effects	151
8.5.1	FUCCI	152
8.5.2	Analysis of the experimental data	152
8.5.3	Expression of the transition rates without circadian control	153
8.5.4	Model simulations	156
8.6	A numerical optimization problem with toxicity constraints	160
8.6.1	Optimization of the Floquet eigenvalue	161
8.6.2	Determination of the gradient of the objective function	163
8.6.3	Simulations	166
8.7	Discussion and future prospects for this work	167
9	Conclusion	171
A	An upper bound on the spectral gap of nonnegative matrices	173
A.1	Reversibilization for nonnegative matrices	173
A.2	Evaluating the spectral gap of a nonnegative matrix	174

Remerciements

Je remercie d'abord mon directeur de thèse Stéphane Gaubert. Sa passion des mathématiques ressort jusque dans son langage de tous les jours. Merci de m'avoir montré comment fonctionne le monde de la recherche en mathématiques. Merci pour toutes tes astuces et pour ta patience.

Merci Marianne Akian de m'avoir aidé à clarifier mon propos lorsque j'écrivais des articles difficiles à comprendre. Ta bonne humeur permanente m'a donné du courage pour débloquent les preuves difficiles. Merci Mustapha Bouhtou pour la confiance que tu m'as donnée et sans laquelle cette thèse n'aurait pas été la même. Il y a bien sûr le financement de la thèse par Orange Labs mais surtout tu as su me présenter clairement les enjeux industriels de mes recherches. Merci aussi d'avoir pu trouver du temps pour suivre et t'intéresser à mes travaux et résultats scientifiques.

Je remercie Jean Clairambault pour notre collaboration fructueuse. Tu m'as permis d'élargir mes connaissances en m'initiant aux mathématiques pour la médecine et j'ai beaucoup apprécié. Merci aussi Frédérique Billy pour tout le travail que tu as fait sur nos articles en commun, en un temps record.

Je remercie mes rapporteurs, Paul Van Dooren et Konstantin Avratchenkov. En plus du travail énorme que vous avez dû fournir pour faire votre rapport sur cette thèse dense et pas toujours parfaitement rédigée, vous m'avez aidé dans ma recherche d'un post-doc pour l'année à venir. Merci à vous Professeurs Tempo et De Lara: vos travaux ont inspiré des idées majeures de ma thèse.

Je remercie aussi la communauté mathématique, ceux qui sont venus à mes exposés et avec lesquels j'ai communiqué. Nos échanges ont mis en évidence ma transformation progressive du statut d'étudiant à celui de docteur. Merci aux rapporteurs de mes articles qui ont proposé des améliorations substantielles.

Merci à tous mes collègues du CMAP, docteurs, doctorants, ingénieurs. Vous m'avez accueilli et si j'ai pu faire la distance entre chez moi et le laboratoire pendant trois ans, c'est un peu parce que je savais que vous seriez là quand j'arriverai. Merci aux personnels d'appui à la recherche, et plus particulièrement à Wallis Filippi pour l'aide inestimable que vous m'avez apportée.

Merci à mes anciens professeurs de l'ENSTA-Paristech, Pierre Carpentier, Hasnaa Zidani et Patrick Ciarlet, de m'avoir permis de devenir un de vos collègues pendant les séances de

travaux dirigés et de travaux pratiques.

Merci à ma famille et à mes amis. Vous avez toujours été là pour me soutenir. Et enfin, merci à toi, Marion. Ton soutien continu m'a fait chaud au coeur. Ton aide m'a été précieuse pour la mise en page de mes articles, la préparation de mes exposés, ton coaching pour les conférences et même de ton aide pour trouver les erreurs cachées dans mes calculs. Merci pour tout et merci encore.

CHAPTER *1*

Introduction

1.1 Context of the thesis*

1.1.1 Web ranking

Web search engines like Google, Yahoo! Search, Bing, Baidu (China) or Voila (France) are unavoidable for any search of information on the World Wide Web. Search engines date back from the 90's when the web began to grow too much for the information to be recovered by browsing only [Nin08].

A search engine is mainly composed of an index of crawled pages and a relevance algorithm that selects and sorts the pages to list to the user. The index is built with the data gathered from computer programs called automatic indexers, crawlers, spiders or bots, that read web pages, select the valuable information and go to another web page by following a hyperlink (or any other crawling rule). A characteristic of web search is the huge size of the index tables [BYRN99]. Google for instance reports to have crawled more than 1 trillion unique urls [AH08]¹. Then, an inverted index is constructed, listing all the web pages associated to each keyword.

* This work was performed as part of a collaboration between INRIA, CMAP, Ecole Polytechnique (Maxplus team) and Orange Labs (Mustapha Bouhtou). The author was supported by Orange Labs through the research contract CRE 3795 with INRIA.

¹On May 19th, 2012, for a search for the keyword “a”, www.google.com reported to have about 25,270,000,000 results, which is an estimate of the number of pages with unique content.

Relevance algorithms consider various features of the pages associated to a query in order to print them in a sorted list to the user. Traditional relevance algorithms are based on the boolean and statistical models [MRS08]. In the boolean model, a query is seen as a boolean formula. For instance for the query “optimization AND eigenvector”, the algorithm will return all the pages that are at the same time in the rows of the inverted index corresponding to the keyword “optimization” and to the keyword “eigenvector”. It is a simple algorithm that selects a subset of the pages but it does not rank the results any further. A possibility is to use statistical information, for instance in the form of term frequencies or the position of the term within the page, to determine the relevance of each document with respect to the query. The vector space model [SM83] is a statistical model that represents the documents and queries as vectors in a multidimensional space, whose dimensions are the terms used to build an index to represent the documents. One then defines the similarity between a document and a query as the cosine of the angle between the vectors representing them. The results are then ranked according to the similarity between the pages and the query.

Early search engines worked with these traditional relevance algorithms. But everybody can write and publish on the web, which implies that many documents are of poor quality, although they may be relevant from a statistical point of view [BYRN99]. Fortunately, the web is organized by the hyperlinks between pages. A hyperlink can be interpreted as a citation of the pointed web page by the pointing page. A popular measure of relevance for a page, that can be combined with classical considerations on the keywords, is the number of hyperlinks that point to that page. In this thesis, we study mathematical problems associated to such link-based ranking algorithms.

Indeed, hyperlink counting quickly became obsolete because web spammers created thousands of artificial pages pointing to the page they wanted to promote [GGM05b]. In 1998, two more sophisticated link-based ranking algorithms were developed: Brin and Page’s PageRank and Kleinberg’s HITS. The PageRank [BP98] is the core of the search engine Google. This innovation has led Google to be the leader in the search engine market and nowadays PageRank is still the reference for link-based web ranking algorithms. It relies on the notion of web graph, which is a digraph with a node for each web page and an arc between pages i and j if there is a hyperlink from page i to page j . The PageRank is defined as the invariant measure of a walk made by a random surfer on the web graph. When reading a given page, the surfer either selects a link from the current page (with a uniform probability), and moves to the page pointed by that link, or interrupts his current search, and then moves to an arbitrary page, which is selected according to given “zapping” probabilities. The rank of a page is defined as its frequency of visit by the random surfer. It is interpreted as the “popularity” of the page.

The HITS algorithm [Kle99] is composed of two steps and the output depends on the query of the user. Given a query, it first selects a seed of pages that are relevant to the query according to their text content. This seed is then extended with pages linking to them, pages to which they link and all the hyperlinks between the pages selected. We thus obtain a subgraph of the web graph focused on the query. Then, the second step assigns each page two scores: a hub score v and an authority score u such that good hubs should point to good authorities and good authorities should be pointed to by good hubs. Introducing the adjacency matrix A of the focused graph, this can be written as $v = \rho Au$ and $u = \rho A^T v$ with $\rho \in \mathbb{R}_+$, which means that the vector of hub scores is the Perron eigenvector of the matrix $A^T A$ and that the vector of authority scores is the Perron eigenvector of AA^T .

Both algorithms rank web pages according to the value of the coordinate of the Perron vector of a given nonnegative matrix. Indeed, by the Perron-Frobenius theorem [BP94], we

know that the spectral radius of a nonnegative matrix is an eigenvalue (called the Perron root) and that there exists an associated eigenvector with nonnegative coordinates. Other link-based ranking algorithms have been developed like SALSA (Lempel and Moran [LM00]), HOTS (Tomlin [Tom03]) or Sinkhorning (Smith [Smi05] and Knight [Kni08] independently). Even if some of them correspond to nonlinear operators, they still rank web pages according to an “eigenvector” of a nonlinear map. The non-linear Perron-Frobenius theory, that deals with monotone and homogeneous maps, is a unifying framework for all those rankings.

1.1.2 Age-structured population dynamics

Tissue proliferation relies on the cell division cycle: one cell becomes two after a sequence of molecular events that are physiologically controlled at each step of the cycle at so-called checkpoints, in particular at transitions between phases of the cycle [Mor06]. Tissue proliferation is the main physiological process occurring in development and later in maintaining the permanence of the organism in adults, at that late stage mainly in fast renewing tissues such as bone marrow, gut and skin.

Proliferation is normally controlled in such a way that tissue homeostasis is preserved. By tissue homeostasis we mean permanence in the mean of tissue in volume, mass and function to ensure satisfaction of the needs of the whole organism. In cancer tissues, this physiological control, which also relies on the so-called checkpoints in the division cycle of individual replicating cells, is disrupted, leading to an overproduction of cells that eventually results in the development of tumours.

Anticancer drugs all attack the cell division cycle, either by slowing it down (possibly until quiescence, i.e., non proliferation, cells remaining alive), or by blocking it at checkpoints, which in the absence of cell material repair eventually leads to cell death.

Various mathematical models [Lai64, Mur90, Web92, PA95, BCL06, LS07a] have been proposed to describe the action of anticancer drugs in order to optimise it, that is to minimise the number of cancer cells or a related quantity, as the growth rate of the cancer cell population. The constraints at stake, met everyday in the clinic of cancers, are related mainly to resistance to treatment in cancer cell populations and to unwanted toxicity in healthy tissues.

The representation of the dynamics of the division cycle in proliferating cell by physiologically structured partial differential equations (PDEs), which dates back to McKendric [McK26], is a natural frame to model proliferation in cell populations, healthy or cancerous. Considering time-periodic parameters, the growth rate of the cell population is modelled by the Floquet eigenvalue of the system of PDEs. In the associated discretized problem, the Floquet eigenvalue is approximated by the Perron eigenvalue of a nonnegative matrix [MMP05]. The Perron vector gives the asymptotic age-distribution of the cells in the discretized problem.

Nonnegative matrices can also be used directly, without a discretization step to model population dynamics. The matrix is then called the Leslie matrix of the biological or ecological system [Cas06]. The Perron eigenvalue is here also interpreted as the growth rate of the population. Hence, although the application is quite different from web ranking, we can see that the main mathematical tool is the same: Perron eigenvalues and Perron vectors of nonnegative matrices.

1.2 Motivation

The web is nowadays a huge marketplace where digital visibility is a question of life or death for many companies. They have massively recourse to advertising in order to attract web surfers to their web site and to fidelize them. Web search engines play a key role in the construction of digital visibility. They sell sponsored links that are printed when an associated query is formulated, but also being well ranked in so-called natural results is an important issue. The importance of optimizing the PageRank, specially for e-business purposes, has even led to the development of a number of companies offering Search Engine Optimization services [CLF09].

The original motivation of this thesis work, carried out as part of a cooperation between INRIA and Orange Labs, is the optimization of the ranking of the pages of a given web site. In practice, the border between a clever website construction and keyword selection and link spamming is tight: a page is considered to be a spam page if it is built more for search engines than for web surfers. We thus pay a particular attention to design constraints that should guarantee that the web site keeps its initial purpose even after the optimization procedure. We shall restrict to link-based rankings and we will not consider keyword selection procedures. We assume that the webmaster controls a given set of hyperlinks and that she wants to get the best rank possible while respecting the design constraints she has fixed. The web ranking optimization problem we consider consists in finding an optimal hyperlink strategy maximizing a scalar function of a given ranking subject to design constraints. PageRank, HITS and SALSA are Perron vector rankings, which means that they correspond to the Perron vector of an elementwise nonnegative matrix. The PageRank optimization problem has been studied in several works [AL06, MV06, dKNvD08, IT09, CJB10] but other web ranking optimization problems have been comparatively less studied.

When optimizing the ranking, we thus optimize a scalar utility function of the Perron eigenvector over a set of nonnegative irreducible matrices. The matrix is constructed from the web graph, so controlling the hyperlinks corresponds to controlling the matrix itself.

Perron eigenvalue and eigenvector optimization does not restrict to web ranking optimization. For instance, Perron eigenvalue optimization has been considered for the signal to interference ratio balancing problem [BS07, TFL11]. In the context of biological and ecological systems, the modelization of the growth rate of a population by the Perron eigenvalue of a nonnegative matrix is a natural framework for its optimization [Men76]. McNamara and Forshlung [MF96] and Logofet [Log08] assume that living beings tend to maximize the growth rate of their species. Hence they deduce behavioral laws or identify hidden parameters. De Lara and Doyen [DLL08] study sustainable management of natural resources via optimal control and Leslie matrices.

For cancer chemotherapeutics, age-structured PDE models have been extensively studied in the last 25 years [Web90, AK93, KKA06, HWAW07]. Then, one searches for drug infusion schedules that minimize the growth rate of cancer cell population while maintaining the growth rate of the healthy cell population over a given toxicity threshold [BCL06].

One way to optimize pharmacological treatments in cancer, taking into account the cell division cycle on which tissue proliferation relies, is to take advantage of the control that circadian clocks are known to exert on it. Such treatments are termed *chronotherapies of cancer* [L02]. The circadian clocks exert a rhythmic regulating control with a period of approximately 24 hours (hence their name: *circa diem* = about one day) on hormonal, metabolic, behavioral and proliferative processes [RW02, SSC02]. It has been observed in clinical settings that patients with cancer, whose circadian rhythms (rest/activity, blood cortisol) were

damped or ablated, showed clinically more fatigue and had poorer life prognosis [KYS⁺01].

In modern oncology, much before molecular circadian clocks were known, it has been experimentally observed that for anticancer drugs, the time of delivery in the 24 h span matters. This has led teams of oncologists to use in the clinic of cancers drug delivery schedules with sinusoidal shape implemented in programmable pumps, with 24 h period and peak delivery times that were determined for each drug as the best possible from trials and errors on laboratory animals and subsequently transferred to the clinic [L02, L06, L08, LACG08, LOD⁺10, LS07b, MWB⁺00].

Altinok et al. developed a cellular automaton for the cell cycle [ALG07a] and showed that an increased variability in the answer to the circadian clock results in less synchronized divisions in the cell population and an increased growth rate. The authors then show that chemotherapy infusions that take into account this desynchronization of cancer cells are more efficient. In [BCF⁺11b], we studied the synchronization of cells in age-structured PDE models and its consequences on cancer chemotherapeutics. When modelling the healthy and cancer cell populations by age-structured PDE models, we are facing a Floquet eigenvalue optimization problem with 24h-periodic parameters and controls, respectively modelling the circadian clock and the drug infusions. The Floquet eigenvalue optimization problem is then approximated by a Perron eigenvalue optimization problem thanks to a discretization step.

The question of the optimization of the eigenvalues has been particularly studied in the case of symmetric matrices, for applications in shape optimization [DK92], composite materials [CL96], optimal design of experiments [Puk06] and many other fields. As the largest eigenvalue defines a convex function of the coefficients of a symmetric matrix, these problems lead most of the time to convex programming, or even semi-definite programming [CDW75, Ove91, SF95, LO96].

The applications that we deal with in this thesis (web ranking, population dynamics) lead to another type of problems, namely the optimization of the Perron value or the Perron vector of a nonnegative, but not necessarily symmetric, matrix depending on parameters. This leads most of the time to non-convex problems, although some convexity properties are sometimes present in an indirect form [Kin61] or in particular cases: Markov decision processes, stochastic matrices [AHH⁺09], independently controlled rows [BN10, NP11].

1.3 Contribution

We first present previous results for PageRank optimization and some link-based web ranking algorithms. We show that most of the link-based ranking algorithms considered in the literature belong to the class of nonlinear Perron vector rankings. This theory gives unified conditions for the rankings to be well-defined, that is existence, uniqueness and positivity of the scores. Thanks to nonlinear Perron-Frobenius theory, we recover the convergence of the power method for these ranking methods and we prove the convergence of several versions of Tomlin's HOTS algorithm (Chapter 6).

In Chapter 3, we prove a general result of independent interest concerning Markov decision processes with implicitly defined action sets. We introduce the notion of *well-described* Markov decision processes, in which, although there may be an *exponential* number of actions, there is a polynomial time *strong separation oracle* for the actions polytope (whereas the classical complexity results assume that the actions are explicitly enumerated [PT87]). We prove, as an application of the theory of Khachiyan's ellipsoid method, that the ergodic control

problem for well-described Markov decision process is polynomial time solvable, even in the multi-chain framework. We also generalize the polynomial time solvability result to Markov decision processes with finite state space and convex action spaces. We give a convex program, the solution of which gives the expected average cost and a bias vector of the problem.

Then, in Chapter 4, we study a continuous and a discrete version of the PageRank optimization problem. We show that the continuous PageRank optimization problem with design constraints defined by the equations of facets of the polytope of admissible transition probabilities of web surfers reduces to an infinite horizon expected average cost problem on a well described Markov Decision Process. Then we show that the discrete PageRank optimization problem, in which in each page, there are obligatory links, facultative links and forbidden links and the transitions are uniform on the hyperlinks present on the page, is equivalent to a relaxed continuous PageRank optimization problem with a concise representation of the action sets. Hence, both problems are solvable in polynomial time. We provide a very efficient algorithm to solve the optimization problem: indeed, we show that optimizing the PageRank is not essentially more difficult than computing it. We study the shape of an optimized web site and we show that there exists a quantity called the mean reward before teleportation which gives a total order of preference in pointing to a page or another. We also extend the ergodic control formulation to problems with constraints coupling the behaviors of several pages. We report numerical results on the web site of one of my co-authors as well as on a fragment of the web (4.10^5 pages from the universities of New Zealand).

In Chapter 5, we give a new link spam detection and PageRank demotion algorithm called MaxRank. Like TrustRank [GGMP04] and AntiTrustRank [KR06], it starts with a seed of hand-picked trusted and spam pages. We define the MaxRank of a page as the frequency of visit of this page by a random surfer minimizing an average cost per time unit. On a given page, the random surfer selects a set of hyperlinks and clicks with uniform probability on any of these hyperlinks. The cost function penalizes spam pages and hyperlink removals. The goal is to determine a hyperlink deletion policy that minimizes this score. The MaxRank is interpreted as a modified PageRank vector, used to sort web pages instead of the usual PageRank vector. We show that the bias vector of the associated ergodic control problem, which is unique up to an additive constant, is a measure of the “spamacity” of each page, used to detect spam pages. We give a scalable algorithm for MaxRank computation that allowed us to perform numerical experiments on the WEBSpAM-UK2007 dataset [web07]. We show that our algorithm outperforms both TrustRank and AntiTrustRank for spam and nonspam page detection.

Then, we study general Perron value and Perron vector optimization problems, in order to apply it to the other web page rankings. Markov decision process techniques do not apply any more but under natural assumptions, the theory of differentiable optimization [BGLS06, Ber95, NW99] applies. As for PageRank optimization, the size of the web graph requires scalable algorithms. We thus concentrate on first order methods, more efficient than second order methods for large problems, and we give a fast algorithm for the computation of the gradient.

We begin this study of web ranking algorithms different from PageRank with proving the convergence of Tomlin’s HOTS algorithm [Tom03] at a linear rate of convergence in Chapter 6. The HOTS vector is the vector of the exponentials of the dual variables of an optimal flow problem. The flow represents an optimal distribution of web surfers on the web graph in the sense of entropy maximization. The dual variable, one by page, is interpreted as the “temperature” of the page. We first study a simplified version of the algorithm, which is a

fixed point scaling algorithm designed to solve the matrix balancing problem for nonnegative irreducible matrices. The proof of convergence is general (nonlinear Perron-Frobenius theory) and applies to a family of deformations of HOTS. Then, we address the effective HOTS algorithm which is the version designed by Tomlin for the ranking of web pages. The model is a network entropy maximization problem which generalizes matrix balancing. We show that, under mild assumptions, the HOTS algorithm converges at a linear rate of convergence. The proof relies on a uniqueness property of the fixed point and on the existence of a Lyapunov function. We also show that the coordinate descent algorithm [SZ90, LT92] can be used to find the ideal and effective HOTS vectors and we compare HOTS and coordinate descent algorithms on fragments of the web graph. Finally, we give an algorithm to compute the HOTS vector when bounds on the flow of websurfers are known and a normalized HOTS algorithm with better convergence properties.

In Chapter 7, we study the problem of optimizing the Perron eigenvector of a controlled matrix and apply it web ranking optimization. Our first main result is the development of a scalable algorithm for the local optimization of a scalar function of the Perron eigenvector over a set of nonnegative irreducible matrices. Indeed, we show that the global Perron vector optimization over a convex set of nonnegative matrices is NP-hard, so we focus on the searching of local optima. We give a power-type algorithm for the computation of the matrix of the partial derivatives of the objective, based on the fact that it is a rank 1 matrix. It shows that computing the partial derivatives of the objective has the same computational cost as computing the Perron vector by the power method, which is the usual method when dealing with the large and sparse matrices built from the web graph. Then we give an optimization algorithm that couples power and gradient iterations. Each step of the optimization algorithm involves a suitable number of power iterations and a descent step. By considering this algorithm to be an approximate projected gradient algorithm [Pol97, PP02], we prove that the algorithm converges to a stationary point. Compared with the case when the number of power iterations is not adapted dynamically, we got a speedup between 3 and 20 in our numerical experiments together with a more precise convergence result.

Then we apply Perron vector optimization to the optimization of scalar functions of HITS authority or HOTS score. We derive optimization algorithms and, thanks to the low rank of the matrix of partial derivatives, we show that the optimal linkage strategies of both problems satisfy a threshold property. This property was already proved for PageRank optimization in [dKNvD08]. As in [IT09, CJB10, FABG13], we partition the set of potential links (i, j) into three subsets, consisting respectively of the set of *obligatory links*, the set of *prohibited links* and the set of *facultative links*. When choosing a subset of the facultative links, we get a graph from which we get any of the three ranking vectors. We are then looking for the subset of facultative links that maximizes a given utility function. We also study the associated relaxed problems, where we accept weighted adjacency matrices. This assumes that the webmaster can influence the importance of the hyperlinks of the pages she controls, for instance by choosing the size of the font, the color or the position of the link within the page. In fact, we shall solve the relaxed problems and then give conditions or heuristics to get an admissible strategy for the discrete problems.

Perron eigenvalues and eigenvectors are useful in a wider framework than web ranking. In the last chapter of this thesis, we present and analyse a mathematical model for the optimization of cancer drug treatments in cycling cell population models with age structure. We consider a drug, 5-FluoroUracil (5-FU), that prevents cells from starting mitosis. The proliferating healthy and cancer cell populations are represented by the same age-structured model of

the McKendrick type, with different physiological controls for the two cases. Each dynamics is given by physiologically structured PDEs where the dynamic variables are the number of cells of each age in each phase of the proliferation cycle: first growth phase (G1), DNA synthesis followed by second growth phase (S-G2) and mitosis (M), and the only structure variable is age in the phase. The parameters of the model are the death rates and transitions from a phase to the next in the cell cycle. In this work, we assume that proliferation is physiologically controlled by a circadian clock, which implies that the coefficients of the model are not only age and phase-dependent but also 24-h periodic functions. Our fundamental hypothesis is that healthy and cancer cells proliferate following the same model, but that cancer cells are characterized by a looser response to the circadian control, which gives them a proliferation advantage. We show how recent fluorescence-based image modelling techniques performed at the single cell level in proliferating cell populations allow one to identify some key parameters of the population dynamics. Then, we consider time-dependent 5-FU infusions that disrupt the transition from phase G2 to M. We study the problem of minimizing the growth rate of the cancer cell population, modeled by the Floquet eigenvalue of the population dynamics, with the constraint that the growth rate of the healthy cell population remains over a given toxicity threshold. The goal is to find (periodic) chemotherapy schedules that are viable in the long term and effective in the fight against cancer. When we discretize the problem, the Floquet eigenvalues are approached by the Perron eigenvalues of sparse nonnegative matrices. We developed a multiplier's method for the local optimization of the growth rates, that takes advantage of a low rank property of the gradient of the Perron eigenvalue. The eigenvalue optimization algorithm is based on the algorithm developed in Chapter 7. We calculated the gradient of the objective function at stake. We implemented the multiplier's method to solve the problem, where the internal unconstrained optimization problems are solved by the coupling of power and gradient iterations.

1.4 Organization

In Chapter 2, we present previous results for PageRank optimization and some link-based web ranking algorithms considered in the literature.

In Chapter 3, we give new results on effective resolution of well-described Markov Decision Processes problems with a possibly exponential number of actions, or convex action spaces. Theorem 3.1 has been published in [FABG13].

In Chapter 4, we show that the PageRank optimization problem can be solved in polynomial time by reducing it to an infinite horizon expected average cost problem on a well described Markov Decision Process. We give an very efficient algorithm to solve the optimization problem: we show that optimizing the PageRank is not essentially more difficult than computing it. Then we also deal with constraints that couple the behavior of several pages. This chapter follows the lines of [FABG13].

In Chapter 5, we base on our results on PageRank optimization to develop a new ranking algorithm, called MaxRank, designed to fight spam pages.

In Chapter 6, we study the convergence properties of Tomlin's HOTS algorithm. Those results are under review in [Fer12a].

In Chapter 7, we study the Perron value and Perron vector optimization problems. We give a scalable algorithm for the local minimization of a scalar function of the Perron vector that uses a low rank property of the matrix of partial derivatives of the criterion and

couples gradient and power iterations. We prove convergence by considering it as an approximate gradient method. We then apply these result to Kleinberg's HITS and Tomlin's HOTS optimization. Those results are under review in [Fer12b].

In Chapter 8, we present another application of Perron value optimization to cancer chemotherapeutics. This work has been published in [BCF⁺11a, BCF⁺11b, BCF12].

Introduction (en français)

1.5 Contexte de la thèse*

1.5.1 Référencement web

Les moteurs de recherches comme Google, Yahoo! Search, Bing, Baidu (Chine) ou Voila (France) sont incontournables pour toute recherche d'information sur le web. L'origine des moteurs de recherches se situe dans les années 1990 où le web commença à devenir trop gros pour que l'information ne soit trouvée uniquement en naviguant sur la toile [Nin08].

Un moteur de recherche est composé principalement d'un index de pages et d'un algorithme de classement qui sélectionne les pages pertinentes et les classe dans une liste retournée à l'utilisateur. L'index est construit à partir des données rassemblées par des programmes informatiques appelés robots d'indexation ou crawlers, qui lisent une page web, sélectionnent les informations importantes et se dirigent vers une autre page en suivant un hyperlien (ou n'importe quelle règle de parcours). Une caractéristique de la recherche sur le web est la grande taille des index [BYRN99]. Par exemple, Google rapporte avoir indexé plus de mille milliards d'urls uniques [AH08]¹. Ensuite, un index inversé est créé. Il liste toutes les pages associées à chaque mot clé.

Les algorithmes de classement considèrent plusieurs caractéristiques des pages associées à une requête avant de les afficher à l'utilisateur dans une liste ordonnée. Les algorithmes de classement traditionnels s'appuient sur le modèle booléen et le modèle statistique [MRS08]. Dans le modèle booléen, une requête est vue comme une formule booléenne. Par exemple, pour la requête "optimisation ET vecteur propre", l'algorithme va retourner toutes les pages qui sont à la fois dans les lignes de l'index inversé correspondant au mot clé "optimisation" et au mot clé "vecteur propre". C'est un algorithme simple qui sélectionne un sous ensemble des pages mais il ne classe pas les résultats obtenus. Une possibilité est alors d'utiliser des

* Ce travail s'inscrit dans le cadre d'une collaboration entre INRIA, CMAP École Polytechnique (équipe-projet Maxplus) et Orange Labs (Mustapha Bouhtou). Son auteur a été soutenu par Orange Labs via le contrat de recherche CRE 3795 avec INRIA.

¹Le 19 mai 2012, pour une recherche avec comme mot clé "a", www.google.com disait avoir environ 25,270,000,000 résultats, ce qui donne une estimation du nombre de pages avec un contenu unique.

statistiques, comme par exemple la fréquence de chaque terme et sa position dans la page, pour déterminer la pertinence de chaque document par rapport à la requête. Le modèle vectoriel [SM83] est un modèle statistique qui représente les documents et les requêtes comme des vecteurs dans un espace multidimensionnel dont les dimensions sont les termes utilisés pour construire l'index qui représente les documents. Ensuite, on définit la similarité entre un document et une requête comme étant le cosinus de l'angle entre les vecteurs les représentant. Les résultats sont ensuite ordonnés en fonction de la similarité entre les pages et la requête.

Au début, les moteurs de recherche fonctionnaient avec ces algorithmes de classement traditionnels. Mais tout le monde peut écrire et publier sur le web, ce qui implique que beaucoup de documents sont de mauvaise qualité, bien qu'ils puissent être pertinents d'un point de vue statistique [BYRN99]. Heureusement, le web est organisé par les hyperliens entre les pages. Un hyperlien peut être interprété comme une citation de la page pointée par la page qui pointe. Une mesure répandue de la pertinence d'une page, qui peut être combinée avec les considérations classiques sur les mots clés, est le nombre d'hyperliens qui pointent vers cette page. Dans cette thèse, nous étudions des problèmes mathématiques associés à ces algorithmes de classement se basant sur les hyperliens.

En fait, le décompte des hyperliens devint rapidement obsolète parce que des spammeurs du web créèrent des milliers de pages artificielles pointant vers la page qu'ils voulaient promouvoir [GGM05b]. En 1998, deux algorithmes plus sophistiqués furent développés : le PageRank par Brin et Page et HITS par Kleinberg. Le PageRank [BP98] est le cœur du moteur de recherche Google. Cette innovation a permis à Google de devenir leader sur le marché du moteur de recherche et maintenant le PageRank est toujours la référence pour les algorithmes de classement se basant sur les hyperliens. Il repose sur la notion de graphe du web, qui est un graphe orienté avec un nœud pour chaque page web et un arc entre les pages i et j si il y a un hyperlien de la page i à la page j . Le PageRank est défini comme la mesure invariante d'une marche faite par un surfeur aléatoire sur le graphe du web. Quand il lit une page donnée, le surfeur peut soit sélectionner un lien sur la page courante (avec une probabilité uniforme) et va sur la page pointée par ce lien, soit interrompre sa recherche et aller à une page arbitraire qui est déterminée en fonction de probabilités de "zapping" données. Le score d'une page est défini comme la fréquence de visite du surfeur aléatoire sur cette page. Il est interprété comme la popularité de la page.

L'algorithme HITS [Kle99] comprend deux étapes et le résultat dépend de la requête de l'utilisateur. Etant donnée une requête, il sélectionne d'abord un ensemble de pages qui sont pertinentes vis à vis de la requête et leur contenu textuel. Cet ensemble de pages est ensuite étendu avec les pages qui pointent vers elles, les pages vers lesquelles elles pointent et tous les hyperliens entre les pages sélectionnées. On obtient alors un sous-graphe du graphe du web centré sur la requête. Ensuite, la deuxième étape assigne à chaque page un score : un score de hub v et un score d'autorité u tels que les bons hubs pointent vers de bonnes autorités et les bonnes autorités sont pointées par de bons hubs. En introduisant la matrice d'adjacence A du graphe centré sur la requête, ceci peut être écrit $v = \rho Au$ et $u = \rho A^T v$ avec $\rho \in \mathbb{R}_+$, ce qui signifie que le vecteur des scores de hub est le vecteur propre de Perron de la matrice $A^T A$ et que le vecteur des scores d'autorité est le vecteur propre de Perron de AA^T .

Les deux algorithmes classent les pages en fonction de la valeur de la coordonnée du vecteur de Perron d'une matrice positive. En fait par le théorème de Perron-Frobenius [BP94], nous savons que le rayon spectral d'une matrice positive terme à terme est une valeur propre (appelée valeur de Perron) et qu'il existe un vecteur propre associé avec des coordonnées positives. D'autres algorithmes de classement utilisant les hyperliens ont été développés comme

SALSA (Lempel et Moran [LM00]), HOTS (Tomlin [Tom03]) ou Sinkhorning (Smith [Smi05] et Knight [Kni08] indépendamment). Même si certains correspondent à des opérateurs non linéaires, ils classent toujours les pages en fonction du “vecteur propre” d’une application non linéaire. La théorie de Perron-Frobenius non linéaire, qui traite des applications monotones et homogènes, est un cadre unificateur pour toutes ces méthodes.

1.5.2 Dynamiques de populations structurées en âge

La prolifération des tissus cellulaires repose sur le cycle de division cellulaire : une cellule se dédouble après une séquence d’évènements moléculaires qui sont contrôlés physiologiquement à chaque étape du cycle au niveau de checkpoints, en particulier au niveau des transitions entre les phases du cycle [Mor06]. La prolifération des tissus est le principal processus physiologique qui se produit dans le développement de l’organisme et plus tard dans son maintien. Chez les adultes, la prolifération a lieu principalement dans les tissus à renouvellement rapide comme la moelle osseuse, les intestins et la peau.

La prolifération est normalement contrôlée de telle manière que l’homéostasie des tissus est préservée. Nous définissons l’homéostasie comme la permanence en moyenne du volume, de la masse et des fonctions des tissus de manière à assurer la satisfaction des besoins de l’organisme. Dans les tissus cancéreux, ce contrôle physiologique, qui repose aussi sur les checkpoints du cycle de division cellulaire de chaque cellule reproductrice, est perturbé, ce qui amène à une surproduction de cellules et finalement développement de tumeurs.

Les médicaments anticancer attaquent tous le cycle de division cellulaire, soit en le ralentissant (parfois jusqu’à la quiescence, c’est-à-dire l’arrêt de la prolifération, les cellules restant vivantes), soit en bloquant le cycle à des checkpoints ce qui, en l’absence de réparation du matériel cellulaire, mène à la mort cellulaire.

Différents modèles mathématiques [Lai64, Mur90, Web92, PA95, BCL06, LS07a] ont proposé de décrire l’action des médicaments anticancer pour l’optimiser, c’est-à-dire pour minimiser le nombre de cellules cancéreuses ou une quantité reliée comme le taux de croissance de la population de cellules cancéreuses. Les contraintes en jeu, rencontrées tous les jours dans le traitement des cancers, sont principalement en lien avec la résistance des populations de cellules cancéreuses au traitement et avec la toxicité indésirable du traitement pour les tissus sains.

La représentation de la dynamique du cycle de division des cellules proliférantes par des équations aux dérivées partielles (EDP) structurées en âge, qui remonte à McKendrick [McK26], est un cadre naturel pour modéliser la prolifération dans les populations de cellules saines ou tumorales. Si on considère des paramètres périodiques par rapport au temps, le taux de croissance de la population de cellules est modélisé par la valeur propre de Floquet du système d’EDP. Dans le problème discrétisé associé, la valeur propre de Floquet est approchée par la valeur propre de Perron d’une matrice positive terme à terme [MMP05]. Le vecteur de Perron donne la distribution asymptotique des âges des cellules dans le problème discrétisé.

Les matrices positives peuvent aussi être utilisées directement, sans une étape de discrétisation pour modéliser des dynamiques de population. La matrice est alors appelée la matrice de Leslie du système biologique ou écologique [Cas06]. La valeur propre de Perron est ici aussi interprétée comme le taux de croissance de la population. Ainsi, bien que l’application soit différente du référencement de pages web, nous pouvons voir que l’outil mathématique de base est le même : les valeurs propres et les vecteurs propres de Perron des matrices positives.

1.6 Motivation

Le web est aujourd'hui un gigantesque marché où la visibilité numérique est une question de vie ou de mort pour de nombreuses entreprises. Elles ont massivement recours à la publicité de manière à attirer les internautes vers leur site web et à les fidéliser. Les moteurs de recherche jouent un rôle capital dans la construction de la visibilité numérique. Ils vendent des liens publicitaires qui sont affichés quand une requête associée est formulée, mais le fait d'être bien placé dans les résultats dits naturels est une question importante. L'importance de l'optimisation du PageRank, en particulier pour le commerce électronique, a même conduit au développement d'un certain nombre d'entreprises qui proposent des services d'optimisation du référencement [CLF09].

La motivation originelle de cette thèse, effectuée dans le cadre d'une collaboration entre INRIA et Orange Labs, est l'optimisation du référencement des pages d'un site web donné. En pratique, la frontière entre site intelligemment construit avec des mots clés choisis et du référencement abusif est étroite : une page est considérée comme du spam si elle est faite plus pour les moteurs de recherche que pour les internautes. Nous regarderons donc avec une attention particulière les contraintes de design qui devraient garantir que le site web garde son but initial même après la procédure d'optimisation. Nous nous restreindrons au classement basé sur les hyperliens et nous ne considérons pas les procédures de sélection des mots clés. Nous supposons que le webmaster commande un ensemble d'hyperliens donné et qu'il veut obtenir la meilleure position possible en respectant les contraintes de design qu'il s'est fixées. Le problème d'optimisation du référencement que nous considérons consiste à trouver une stratégie optimale d'hyperliens qui maximise une fonction scalaire d'un classement donné sous des contraintes de design. Le PageRank, HITS et SALSA sont des classements par le vecteur de Perron, ils correspondent donc au vecteur de Perron d'une matrice positive terme à terme. Le problème d'optimisation du PageRank a été étudié dans plusieurs travaux [AL06, MV06, dKNvD08, IT09, CJB10] mais les autres problèmes d'optimisation du référencement ont été comparativement moins étudiés.

Quand on optimise le référencement des pages web, on optimise donc une fonction d'utilité scalaire du vecteur propre de Perron sur un ensemble de matrices positives irréductibles. La matrice est construite à partir du graphe du web donc commander les hyperliens revient à commander la matrice elle-même.

L'optimisation de la valeur propre et du vecteur propre de Perron ne se limite pas à l'optimisation du référencement. Par exemple, l'optimisation de la valeur de Perron a été considérée pour le problème de l'équilibrage du ratio signal sur interférence [BS07, TFL11]. Dans le contexte des systèmes biologiques et écologiques, la modélisation du taux de croissance d'une population par la valeur de Perron d'une matrice positive est un cadre naturel pour son optimisation [Men76]. McNamara et Forshlung [MF96] et Logofet [Log08] font l'hypothèse que les êtres vivants ont tendance à maximiser le taux de croissance de leur espèce. Ils en déduisent ainsi des lois de comportement ou ils identifient des paramètres cachés. De Lara et Doyen [DLL08] étudient la gestion durable des ressources naturelles en utilisant la théorie de la commande optimale et des matrices de Leslie

Pour la chimiothérapie du cancer, les modèles d'EDP structurés en âge ont été largement étudiés dans les dernières 25 années [Web90, AK93, KKA06, HAW07]. On cherche alors des programmes d'injection de médicament qui minimisent le taux de croissance de la population de cellules cancéreuses tout en maintenant le taux de croissance de la population de cellules saines au dessus d'un seuil de toxicité donné [BCL06].

Une manière d’optimiser les traitements pharmacologiques du cancer qui prend en compte le cycle de division cellulaire sur lequel repose la prolifération des tissus est de tirer profit du contrôle exercé sur lui par les horloges circadiennes. De tels traitements sont appelés des *chronothérapies du cancer* [LÓ2]. Les horloges circadiennes exercent un contrôle régulateur rythmé avec une période d’environ 24 heures (d’où leur nom : *circa diem* = environ un jour) sur les processus hormonaux, métaboliques, comportementaux et prolifératifs [RW02, SSC02]. Il a été observé en clinique que les patients atteints du cancer dont les rythmes circadiens (repos/activité, cortisol sanguin) étaient aplatis ou enlevés montraient cliniquement plus de fatigue et avaient des pronostics de vie plus faibles [KYS⁺01].

Dans l’oncologie moderne, bien avant que les horloges circadiennes moléculaires ne soient connues, il a été observé expérimentalement que pour les médicaments anticancer, l’heure de délivrance au cours de la journée a une importance. Ceci a conduit les équipes d’oncologistes à utiliser dans le traitement des cancers des programmes d’injection de médicament avec des formes sinusoïdales, implémentés dans des pompes programmables avec des périodes de 24 h et des pics de distribution qui ont été déterminés au mieux pour chaque médicament par tâtonnements sur des animaux de laboratoire et transférés ensuite en clinique [LÓ2, LÓ6, LÓ8, LACG08, LOD⁺10, LS07b, MWB⁺00].

Altinok et al. ont développé un automate cellulaire pour le cycle cellulaire [ALG07a] et montré qu’une variabilité accrue dans la réponse à l’horloge circadienne induit des divisions moins synchronisées dans la population de cellules et un taux de croissance accru. Les auteurs ont ensuite montré que les injections de chimiothérapie qui prennent en compte cette désynchronisation des cellules cancéreuses sont plus efficaces. Dans [BCF⁺11b], nous avons étudié la synchronisation des cellules dans in modèle d’EDP structurées en âge et ses conséquences sur la chimiothérapeutique du cancer. Quand on modélise les populations de cellules saines et cancéreuses par des EDP structurées en âge, nous sommes confrontés un problème d’optimisation de la valeur propre de Floquet d’un système avec des paramètres et une commande périodiques de période 24 h, qui représentent respectivement l’horloge circadienne et les injections de médicament. Le problème d’optimisation de la valeur propre de Floquet en ensuite approché par un problème d’optimisation de la valeur propre de Perron grâce à une étape de discrétisation.

La question de l’optimisation des valeurs propres a été particulièrement étudiée dans le cas des matrices symétriques, pour des applications en optimisation de forme [DK92], matériaux composites [CL96], plans d’expérience optimaux [Puk06] et de nombreux autres domaines. Comme la plus grande valeur propre définit une fonction convexe des coefficients d’une matrice symétrique, ces problèmes se réduisent la plupart du temps à de l’optimisation convexe ou même de l’optimisation semi-définie positive. [CDW75, Ove91, SF95, LO96].

Les application que nous traitons dans cette thèse (référencement web, dynamique de populations) conduisent à un autre type de problèmes : l’optimisation de la valeur de Perron ou du vecteur de Perron d’une matrice positive terme à terme, mais pas nécessairement symétrique, qui dépend de paramètres. Cela donne la plupart du temps des problèmes non convexes, bien que quelques propriétés de convexité soient parfois présentes dans une forme indirecte [Kin61] ou dans des cas particuliers : processus de décision markoviens, matrices stochastiques [AHH⁺09], lignes commandées indépendamment [BN10, NP11].

1.7 Contribution

Nous présentons d’abord les résultats antérieurs sur l’optimisation du PageRank et quelques algorithmes de classement de pages utilisant la structure d’hyperliens du web. Nous montrons que la plupart des algorithmes considéré dans la littérature appartiennent à la classe des classements par vecteur de Perron non linéaire. Cette théorie donne des conditions unifiées pour déterminer si les classements sont bien définis, c’est-à-dire que les scores existent, sont uniques et positifs. Grâce à la théorie de Perron-Frobenius non linéaire, nous retrouvons la convergence de la méthode de la puissance pour ces méthodes de classement et nous prouvons la convergence de plusieurs versions de l’algorithme HOTS de Tomlin (chapitre 6).

Dans le chapitre 3, nous prouvons un résultat général d’intérêt indépendant au sujet des processus de décision markoviens avec des ensembles d’action implicites. Nous introduisons la notion de processus de décision markoviens *bien décrits*, pour lesquels, bien qu’il puisse y avoir un nombre *exponentiel* d’actions, il y a un *oracle polynomial de séparation forte* pour le polytope des actions (alors que les résultats de complexité classiques supposent que les actions sont énumérées explicitement [PT87]). Nous prouvons, comme une application de la théorie de l’ellipsoïde de Khachiyan, que le problème de commande ergodique pour les processus de décision markoviens bien décrits est résoluble en temps polynomial, même dans le cas multichaîne. Nous généralisons aussi ce résultat de résolution en temps polynomial au processus de décision markoviens avec un ensemble d’états fini et des ensembles d’actions convexes. Nous donnons un problème d’optimisation dont la solution donne le coût moyen espéré et un vecteur de biais du problème de décision markovien.

Ensuite, dans le chapitre 4, nous étudions une version continue et une version discrète du problème d’optimisation du PageRank. Nous montrons que le problème continu d’optimisation du PageRank avec des contraintes de design définies par les équations des facettes du polytope des probabilités de transition admissibles du surfeur aléatoire se réduit à un problème de coût moyen espéré en horizon infini sur un processus de décision markovien bien décrit. Nous montrons ensuite que le problème discret d’optimisation du PageRank, dans lequel à chaque page il y a des liens obligatoires, des liens facultatifs et des liens interdits et les transition sont uniformes sur les hyperliens présents sur la page est équivalent à un problème relâché continu d’optimisation du PageRank avec une description concise des ensembles d’actions. Ainsi, les deux problèmes sont résolubles en temps polynomial. Nous donnons un algorithme très efficace pour résoudre le problème d’optimisation du PageRank : en effet, nous montrons qu’optimiser le PageRank n’est pas fondamentalement plus difficile que le calculer. Nous étudions la forme d’un site web optimisé et nous montrons qu’il existe une quantité appelée revenu moyen avant téléportation qui donne un ordre total de préférence sur le fait de pointer vers une page ou une autre. Nous étendons aussi la formulation par la commande ergodique aux problèmes avec des contraintes qui couplent le comportement de plusieurs pages. Nous communiquons des résultats numériques sur le site web d’un de mes co-auteurs et sur un fragment du web de 4.10^5 pages du site web des universités de Nouvelle-Zélande.

Dans le chapitre 5, nous donnons un nouvel algorithme de détection du spam de liens et de rétrogradation du PageRank appelé MaxRank. Comme le TrustRank [GGMP04] et l’AntiTrustRank [KR06], il commence avec un ensemble de pages sélectionnés à la main et étiquetées honnêtes ou spam. Nous définissons le MaxRank d’une page comme la fréquence de visite de cette page par un surfeur aléatoire qui minimise un coût moyen par unité de temps. Sur une page donnée, le surfeur aléatoire choisit un ensemble d’hyperliens et clique avec une probabilité uniforme sur un de ces hyperliens. La fonction coût pénalise les pages

de spam et les suppressions d’hyperliens. Le but est de déterminer un politique de suppression d’hyperliens qui minimise ce score. Le MaxRank est interprété comme un vecteur de PageRank modifié, utilisé pour classer les pages à la place du vecteur de PageRank habituel. Nous montrons que le vecteur de biais du problème de commande ergodique associé, qui est unique à une constante additive près, est une mesure de la “spamité” de chaque page, utilisée pour détecter les pages de spam. Nous donnons un algorithme qui passe à l’échelle pour le calcul du MaxRank qui nous permet d’effectuer des expériences numériques sur la base de données WEBSpAM-UK2007 [web07]. Nous montrons que notre algorithme fait mieux que TrustRank et AntiTrustRank pour la détection des pages de spam et des pages honnêtes.

Ensuite, nous étudions les problèmes généraux d’optimisation de la valeur et du vecteur de Perron, dans le but de l’appliquer à d’autres classements de pages web. Les processus de décision markoviens ne s’appliquent plus mais sous des hypothèses naturelles, la théorie de l’optimisation différentiable [BGLS06, Ber95, NW99] s’applique. Comme pour l’optimisation du PageRank, la taille du graphe du web nécessite des algorithmes qui passent à l’échelle. Nous nous concentrons donc sur les méthodes du premier ordre, plus efficaces que les méthodes du second ordre pour les grands problèmes, et nous donnons un algorithme rapide pour le calcul du gradient.

Nous commençons cette étude des algorithmes de classement différents du PageRank par la preuve que l’algorithme HOTS de Tomlin [Tom03] converge géométriquement dans le chapitre 6. Le vecteur de HOTS est le vecteur des exponentielles des variables duales d’un problème de flot optimal. Le flot représente une distribution optimale d’internautes sur le graphe du web au sens de la maximisation de l’entropie. La variable duale, une par page, est interprétée comme la “température” de la page. Nous étudions d’abord une version simplifiée de l’algorithme, qui est un algorithme de mise à l’échelle conçu pour résoudre le problème d’équilibrage de matrice (matrix balancing) pour les matrices positives irréductibles. La preuve de convergence est générale (théorie de Perron-Frobenius non linéaire) et s’applique à une famille de déformations de HOTS. Ensuite, nous abordons l’algorithme effectif de HOTS qui est la version conçue par Tomlin pour le classement des pages web. Le modèle est un problème de maximisation d’entropie dans un réseau qui généralise l’équilibrage de matrice. Nous montrons que sous des hypothèses modérées, l’algorithme HOTS converge géométriquement. La preuve s’appuie sur une propriété d’unicité du point fixe et sur l’existence d’une fonction de Lyapunov. Nous montrons aussi que la méthode de relaxation (descente le long des coordonnées) [SZ90, LT92] peut être utilisée pour trouver les vecteurs de HOTS idéal et effectif et nous comparons HOTS et la méthode de relaxation sur des fragments du graphe du web. Finalement, nous donnons un algorithme pour calculer le vecteur de HOTS quand des bornes sur le flot d’internautes sont connues et un algorithme de HOTS normalisé avec de meilleures propriétés de convergence.

Dans le chapitre 7, nous étudions le problème de l’optimisation du vecteur propre de Perron d’une matrice commandée et nous l’appliquons à l’optimisation du référencement. Notre premier résultat principal est le développement d’un algorithme qui passe à l’échelle pour l’optimisation locale d’une fonction scalaire du vecteur propre de Perron sur un ensemble de matrices positives irréductibles. En effet, nous montrons que le problème de l’optimisation globale du vecteur de Perron sur un ensemble de matrices positives est NP-difficile, ce qui fait que nous nous concentrons sur la recherche d’optimums locaux. Nous donnons un algorithme de type puissance pour le calcul de la matrice des dérivées partielles de la fonction objectif, en nous appuyant sur le fait que c’est une matrice de rang 1. Cela montre que calculer les dérivées partielles de l’objectif a le même coût en calcul que calculer le vecteur propre de

Perron par la méthode de la puissance, qui est la méthode habituelle quand on considère les matrices grandes et creuses construites à partir du graphe du web. Ensuite, nous donnons un algorithme d'optimisation qui couple les itération puissance et gradient. Chaque étape de l'algorithme fait appel à un nombre adéquat d'itérations puissance et à un pas de descente. En considérant cet algorithme comme un algorithme de gradient projeté approché [Pol97, PP02], nous prouvons qu'il converge vers un point stationnaire. Comparé avec le cas où le nombre d'itérations puissance n'est pas adapté dynamiquement, nous avons obtenu une accélération entre 3 et 20 fois dans nos expériences numériques tout en ayant des résultats de convergence plus précis.

Ensuite, nous appliquons l'optimisation du vecteur de Perron à l'optimisation de fonctions scalaires des scores d'autorité de HITS et de HOTS. Nous dérivons des algorithmes d'optimisation et, grâce à la propriété de petit rang de la matrice des dérivées partielles, nous montrons que les stratégies optimales d'hyperliens des deux problèmes vérifient une propriété de seuil. Cette propriété était déjà prouvée pour l'optimisation du PageRank dans [dKNvD08]. Comme dans [IT09, CJB10, FABG13], nous partitionnons l'ensemble des liens potentiels (i, j) en trois sous-ensembles : les ensembles des liens obligatoires, facultatifs et interdits. Quand on choisit un sous-ensemble des liens facultatifs, on obtient un graphe duquel on peut obtenir les trois vecteurs de classement. Nous cherchons alors le sous-ensemble de liens facultatifs qui maximise une fonction d'utilité donnée. Nous étudions aussi les problèmes relâchés associés où nous acceptons les matrices d'adjacence à poids. Cela suppose que le webmaster peut influencer sur l'importance respective des liens qu'il commande, par exemple en modifiant la taille de la police, la couleur ou la position de l'hyperlien dans la page. En fait, nous résoudrons les problèmes relâchés et nous donnerons ensuite des conditions ou des heuristiques pour obtenir une stratégie admissible pour les problèmes discrets.

Les valeurs propres et les vecteurs propres de Perron sont utiles dans un cadre plus large que le référencement web. Dans le dernier chapitre de cette thèse, nous présentons et nous analysons un modèle mathématique pour l'optimisation des traitements médicamenteux du cancer dans des modèles de populations de cellules proliférantes avec une structure par âges. Nous considérons un médicament, le 5-FluoroUracil (5-FU), qui empêche les cellules de commencer la mitose. Les populations de cellules proliférantes saine et cancéreuse sont représentées par le même modèle structuré en âge du type McKendrick, mais avec des contrôles physiologiques différents dans les deux cas. Chaque dynamique est donnée par des EDP physiologiquement structurées où les variables dynamiques sont les nombres de cellules de chaque âge dans chaque phase : première phase de croissance (G1), synthèse de l'ADN suivie par une deuxième phase de croissance (S-G2) et mitose (M). L'unique variable de structure est l'âge dans la phase. Les paramètres du modèle sont les taux de mort et les taux de transition d'une phase à la suivante dans le cycle cellulaire. Dans ce travail, nous supposons que la prolifération est contrôlée physiologiquement par une horloge circadienne, ce qui implique que les coefficients du modèle dépendent non seulement de l'âge et de la phase mais que ce sont des fonctions 24h-périodiques. Notre hypothèse fondamentale est que les cellules saines et cancéreuses prolifèrent en suivant le même modèle mais que les cellules cancéreuses sont caractérisées par une réponse plus relâchée au contrôle circadien, ce qui leur donne un avantage prolifératif. Nous montrons comment des techniques d'imagerie récentes utilisant la fluorescence et mise en route au niveau individuel des cellules dans des populations de cellules proliférantes nous a permis d'identifier des paramètres clés de la dynamique de population. Ensuite, nous considérons des injections de 5-FU dépendant du temps qui perturbent la transition de la phase G2 à la phase M. Nous étudions le problème de la minimisation

du taux de croissance de la population de cellules cancéreuses, modélisée par la valeur propre de Floquet de la dynamique de population, avec la contrainte que le taux de croissance de la population de cellules saines reste au dessus d'un seuil de toxicité donné. Le but est de trouver des programmes de chimiothérapie périodiques qui sont viables sur le long terme et efficaces dans la lutte contre le cancer. Quand nous discrétisons le problème, les valeurs propres de Floquet sont approchées par les valeurs propres de Perron de matrices positives creuses. Nous avons mis en place une méthode des multiplicateurs pour l'optimisation locale des taux de croissance, qui profite d'une propriété de petit rang du gradient de la valeur propre de Perron. L'algorithme d'optimisation de la valeur propre s'appuie sur l'algorithme développé dans le chapitre 7. Nous avons calculé le gradient de la fonction objectif en jeu. Nous avons implémenté la méthode des multiplicateurs pour résoudre le problème, où les problèmes d'optimisation non contraints internes sont résolus par l'algorithme couplant les itérations puissance et gradient.

1.8 Organisation

Dans le chapitre 2, nous présentons les résultats antérieurs pour l'optimisation du PageRank et quelques algorithmes de classement de pages utilisant la structure d'hyperliens du web.

Dans le chapitre 3, nous donnons des résultats nouveaux sur la résolution effective des problèmes de décision markoviens bien décrits avec un nombre d'actions qui peut être exponentiel ou des espaces d'actions convexes. Le théorème 3.1 a été publié dans [FABG13].

Dans le chapitre 4, nous montrons que le problème d'optimisation du PageRank peut être résolu en temps polynomial en le réduisant à un problème de coût moyen espéré en horizon infini sur un processus de décision markovien bien décrit. Nous donnons un algorithme très efficace pour résoudre le problème d'optimisation : nous montrons qu'optimiser le PageRank n'est pas fondamentalement plus difficile que le calculer. Ensuite, nous traitons des problèmes avec des contraintes qui couplent le comportement de plusieurs pages. Ce chapitre suit les lignes de [FABG13].

Dans le chapitre 5, nous nous appuyons sur nos résultats sur l'optimisation du PageRank pour développer un nouvel algorithme de classement appelé MaxRank fait pour combattre le spam de liens.

Dans le chapitre 6, nous étudions la convergence de l'algorithme HOTS de Tomlin. Ces résultats ont été soumis dans [Fer12a].

Dans le chapitre 7, nous étudions les problèmes d'optimisation de la valeur propre et du vecteur propre de Perron. Nous donnons un algorithme efficace pour le calcul de la matrice des dérivées partielles du critère, qui utilise la propriété de petit rang de cette matrice. Nous donnons un algorithme qui passe à l'échelle qui couple les itérations gradient et puissance et donne un minimum local du problème d'optimisation du vecteur de Perron. Nous prouvons la convergence en le considérant comme une méthode de gradient approché. Nous appliquons ensuite ces résultats à l'optimisation de HITS de Kleinberg et HOTS de Tomlin. Ces résultats ont été soumis dans [Fer12b].

Dans le chapitre 8, nous présentons une autre application de l'optimisation de la valeur propre de Perron à la chimiothérapie. Ce travail a été publié dans [BCF⁺11a, BCF⁺11b, BCF12].

CHAPTER 2

Web ranking and (nonlinear) Perron Frobenius theory

2.1 Google's PageRank

One of the main ranking methods relies on the PageRank introduced by Brin and Page [BP98]. It is defined as the invariant measure of a walk made by a random surfer on the web graph. When reading a given page, the surfer either selects a link from the current page (with a uniform probability), and moves to the page pointed by that link, or interrupts his current search, and then moves to an arbitrary page, which is selected according to given “zapping” probabilities. The rank of a page is defined as its frequency of visit by the random surfer. It is interpreted as the “popularity” of the page.

The PageRank has motivated a number of works, dealing in particular with computational issues. Classically, the PageRank vector is computed by the power algorithm [BP98]. There has been a considerable work on designing new, more efficient approaches for its computation [Ber05, LM06]: Gauss-Seidel method [ANTT02], aggregation/disaggregation [LM06] or distributed randomized algorithms [NP09, IT10, ITB12]. Other active fields are the development of new ranking algorithms [BRR05] or the study of the web graph [BL04].

We recall here the basic elements of the Google PageRank computation. We call *web graph* the directed graph with a node per web page and an arc from page i to page j if page i contains a hyperlink to page j . We identify the set of pages to $[n] := \{1, \dots, n\}$.

Let N_i denote the number of hyperlinks contained in page i . Assume first that $N_i \geq 1$ for

all $i \in [n]$, meaning that every page has at least one outlink. Then, we construct the $n \times n$ stochastic matrix S , which is such that

$$S_{i,j} = \begin{cases} N_i^{-1} & \text{if page } j \text{ is pointed to from page } i \\ 0 & \text{otherwise} \end{cases} \quad (2.1)$$

This is the transition matrix of a Markov chain modeling the behavior of a surfer choosing a link at random, uniformly among the ones included in the current page and moving to the page pointed by this link. The matrix S only depends of the web graph.

We also fix a row vector $z \in \mathbb{R}_+^n$, the *zapping* or *teleportation* vector, which must be stochastic (so, $\sum_{j \in [n]} z_j = 1$), together with a *damping factor* $\alpha \in [0, 1]$ and define the new stochastic matrix

$$P = \alpha S + (1 - \alpha)ez \quad (2.2)$$

where e is the (column) vector in \mathbb{R}^n with all entries equal to 1.

Consider now a Markov chain $(X_t)_{t \geq 0}$ with transition matrix P , so that for all $i, j \in [n]$, $\mathbb{P}(X_{t+1} = j | X_t = i) = P_{i,j}$. Then, X_t represents the position of a websurfer at time t : when at page i , the websurfer continues his current exploration of the web with probability α and moves to the next page by following the links included in page i , as above, or with probability $1 - \alpha$, stops his current exploration and then teleports to page j with probability z_j .

When some page i has no outlink, $N_i = 0$, and so the entries of the i th row of the matrix S cannot be defined according to (2.1). Then, we set $S_{i,j} := z_j$. In other words, when visiting a page without any outlink, the websurfer interrupts its current exploration and teleports to page j again with probability z_j . It is also possible to define another probability vector Z (different from z) for the teleportation from these “dangling nodes”.

The *PageRank* π is defined as the invariant measure of the Markov chain $(X_t)_{t \geq 0}$ representing the behavior of the websurfer. This invariant measure is unique if $\alpha < 1$, or if P is irreducible.

Typically, one takes $\alpha = 0.85$, meaning that at each step, a websurfer interrupts his current search with probability $0.15 \simeq 1/7$. The advantages of the introduction of the damping factor and of the teleportation vector are well known. First, it guarantees that the power algorithm converges to the PageRank with a geometric rate α independent of the size (and other characteristics) of the web graph. In addition, the teleportation vector may be used to “tune” the PageRank if necessary. By default, $z = e^T/n$ is the uniform stochastic vector. We will assume in the sequel that $\alpha < 1$ and $z_j > 0$ for all $j \in [n]$, so that P is irreducible.

The graph on Figure 2.1 represents a fragment of the web graph. We obtained the graph by performing a crawl of our laboratory with 1500 pages. We set the teleportation vector in such a way that the 5 surrounding institutional pages are dominant. The teleportation probabilities to these pages were taken to be proportional to the PageRank (we used the Google Toolbar, which gives a rough indication of the PageRank, on a logarithmic scale). After running the PageRank algorithm on this graph, we found that within the controlled site, the main page of this author has the biggest PageRank (consistently with the results provided by Google search).

2.2 PageRank optimization in the literature

The PageRank optimization problem has been studied in several works: [AL06, MV06, dKNvD08, IT09, CJB10]. In this section, we review quickly the different results presented in




Figure 2.1: The web site of one of my co-authors (colored) and the surrounding sites (white). This 1500-page fragment of the web is aggregated for presentation, using the technique described in [LM06]. The sizes of the circles follow the log of their PageRank.

these works.

Avrachenkov and Litvak analyzed in [AL06] the case of a single controlled page. They introduced $m_{i,j}$, the average time needed to reach j starting from i when the random walk follows the original hyperlink matrix S , i.e., $\alpha = 1$. We refer to the $m_{i,j}$'s as mean first passage times. Using the fact that $m_{i,i} > 1$ verifies $\pi_i = 1/m_{i,i}$, they showed the following:

Proposition 2.1 ([AL06]). *The optimal linking strategy for a Web page is to have only one outgoing link pointing to a Web page with a shortest mean first passage time back to the original page.*

In [MV06], Mathieu and Viennot established several bounds indicating to what extent the rank of the pages of a multi-page website can be changed, and derived an optimal referencing strategy in a special unconstrained case:

Proposition 2.2 ([MV06]). *If the webmaster can fix arbitrarily the hyperlinks in a web site, then, it is optimal to delete every link pointing outside the web site.*

To avoid such degenerate strategies, De Kerchove, Ninove and Van Dooren [dKNvD08] studied the problem of maximizing the sum of the PageRank coordinates in a web site, provided that from each page, there is at least one path consisting of hyperlinks and leading to an external page. They gave a necessary structural condition satisfied by an optimal outlink strategy. Let $([n], \mathcal{E})$ be the web graph and let I be a subset of $[n]$ representing a web site. They denote by e_I the vector such that $e_I(i) = 1$ if $i \in I$ and $e_I(i) = 0$ otherwise. They define $\mathcal{E}(I) = \{(i, j) \in \mathcal{E} \mid i \in I, j \in I\}$, the set of internal links, and

$\mathcal{E}_{\text{out}(I)} = \{(i, j) \in \mathcal{E} \mid i \in I, j \notin I\}$, the set of outgoing links. They introduced the mean number of visits before zapping defined by $v = (1 - \alpha S)^{-1} e_I$. Recall that S is the original transition matrix (without damping factor).

Proposition 2.3 (Theorem 12 in [dKNvD08]). *Let $\mathcal{E} \setminus (\mathcal{E}(I) \cup \mathcal{E}_{\text{out}(I)})$ be given. Let $\mathcal{E}(I)$ and $\mathcal{E}_{\text{out}(I)}$ be such that $\pi \cdot e_I$ is maximal under the assumption that from each page of I , there is at least one path consisting of hyperlinks and leading to an external page. Then there exists a permutation of the indices such that $I = \{1, 2, \dots, n_I\}$,*

$$v_1 > \dots > v_{n_I} > v_{n_I+1} \geq \dots \geq v_n ,$$

and $\mathcal{E}(I)$ and $\mathcal{E}_{\text{out}(I)}$ have the following structure:

$$\begin{aligned} \mathcal{E}(I) &= \{(i, j) \in I \times I \mid j \leq i \text{ or } j = i + 1\} , \\ \mathcal{E}_{\text{out}(I)} &= \{(n_I, n_I + 1)\} . \end{aligned}$$

In [Nin08], Ninove developed a heuristic based on these theoretical results, which was experimentally shown to be efficient.

In [IT09], Ishii and Tempo investigated the sensitivity of the PageRank to fragile (i.e. erroneous or imperfectly known) web data, including fragile links (servers not responding, links to deleted pages, etc.). Instead of focusing on a given website, they proposed more general strategies, where controlled hyperlinks can be any hyperlink, not necessarily the outlinks of a website. They gave bounds on the possible variation of PageRank and introduced an approximate PageRank optimization problem, which they showed to be equivalent to a linear program. They considered the interval matrix $\mathcal{P} = [\underline{P}, \bar{P}]$ such that for all selection of fragile links, the resulting transition matrix $P \in \mathcal{P}$. Then defining $P_c = \frac{1}{2}(\underline{P} + \bar{P})$ and $\Delta = \frac{1}{2}(\underline{P} - \bar{P})$, they introduced the polytope

$$\begin{aligned} \mathcal{Z} = \left\{ z \in \mathbb{R}^{n-1} \mid - (P_c + \Delta - I)Ez \leq (P_c + \Delta - I)g, \right. \\ \left. (P_c - \Delta - I)Ez \leq -(P_c - \Delta - I)g, \sum_{i \in [n]} z_i \leq 1, z \geq 0 \right\} \end{aligned}$$

where $E = \begin{bmatrix} I \\ -1^T \end{bmatrix} \subseteq \mathbb{R}^{n \times (n-1)}$ and $g = [0, \dots, 0, 1]^T \in \mathbb{R}^n$. They then showed that for each choice of fragile links, the vector $[\pi_1, \pi_2, \dots, \pi_{n-1}]^T$ is a vector of \mathcal{Z} . For $l \in [n-1]$, maximizing (resp. minimizing) z_l on \mathcal{Z} thus gives an upper bound (resp. lower bound) on the maximal (resp. minimal) possible value for π_l .

In [CJB10], (see also [CJB09] for more details), Csáji, Jungers and Blondel thought of fragile links as the links that a webmaster controls and studied the PageRank optimization problem in this framework. As the PageRank is the inverse of the mean first return passage time, they reformulated the maximization (resp. minimization) of the PageRank of Page i as the minimization (resp. maximization) of the mean first return passage time to Page i . Thus, they obtain a stochastic shortest path problem. However, the number of actions at a given page is 2^m where m is the number of fragile links on this page. In order to have a polynomial-time algorithm for the optimization of the PageRank of a page, they gave a new problem with an augmented state space. The idea of [CJB10] is that if the control takes place on hyperlinks instead of transition probabilities, there are only two possible actions by facultative link: active or not. They thus proposed the graph augmentation described

in Figure 2.2. If the page i has only facultative hyperlinks, it may have no hyperlink. The state augmentation described before is then not valid any more because having no hyperlinks stops the Markov chain propagation instead of forcing teleportation, as it should be according to the PageRank model. In order to cope with this problem, they proposed an alternative graph augmentation (Figure 2.3). They showed that the original and the graph augmented stochastic shortest path problems are equivalent. Then they showed that the graph augmented stochastic shortest path problems has a polynomial number of states and actions and thus that it is solvable in polynomial time (see Section 3.1 below).




Figure 2.2: The graph transformation of [CJB10]. Left: original Markov chain. Dotted lines represent fragile links while solid links are not fragile. An action consists in choosing a subset of the fragile edges. Right: state augmented Markov chain. The original states of the web graph are represented by squares. One dummy state i' (circle) is added to dispatch websurfers that have not teleported and then all the decisions for the link strategy will take place on the states representing controlled links (diamonds). An action consists in choosing one of the two edges on each state representing a controlled link.




Figure 2.3: Fragile node case. Left: original Markov chain, an action consists in choosing a subset of the facultative links. Right: state augmented Markov chain. An action consists in choosing one of the edges from node i' .

2.3 Other web ranking algorithms

The PageRank is the most widely used link-based web ranking algorithm. However other algorithms have been developed. For instance, some rankings are determined more generally by the Perron eigenvector (i.e. the principal eigenvector) of a nonnegative, but not necessarily stochastic, matrix. The Perron-Frobenius theorem (see [BP94] for instance) states that any nonnegative matrix A has a nonnegative principal eigenvalue called the Perron root and associated nonnegative principal eigenvectors. If, in addition, A is irreducible, then the Perron

root is simple and the (unique up to a multiplicative constant) nonnegative eigenvector, called the Perron vector, has only positive entries. This property makes it a good candidate to sort web pages. The ranking algorithms considered differ in the way of constructing from the web graph a nonnegative irreducible matrix from which we determine the Perron vector. Then, the greater is the Perron vector's coordinate corresponding to a web page, the higher this web page is in the ranking. In [Kee93], such a ranking is proposed for football teams. The paper [Saa87] uses the Perron vector to rank teachers from pairwise comparisons. See also [Vig09b] for a survey on the subject. When it comes to web page ranking, the PageRank is the Perron eigenvector of the transition matrix described above but the HITS and SALSA algorithms also rank pages according to a Perron vector. This class of ranking algorithms motivated our study of Perron vector optimization problems, where we want to optimize a scalar function of the Perron eigenvector over a set of nonnegative irreducible matrices.

The HITS algorithm [Kle99] developed by Kleinberg is not purely a link-based algorithm. It is composed of two steps and the output depends on the query of the user. Given a query, we first select a seed of pages that are relevant to the query according to their text content. This seed is then extended with pages linking to them, pages to which they link and all the hyperlinks between the pages selected. We thus obtain a subgraph of the web graph focused on the query. Then, the second step assigns each page two scores: a hub score v and an authority score u such that good hubs should point to good authorities and good authorities should be pointed to by good hubs. Introducing the adjacency matrix A of the focused graph, this can be written as $v = \rho Au$ and $u = \rho A^T v$ with $\rho \in \mathbb{R}_+$, which means that the vector of hub scores is the Perron eigenvector of the matrix $A^T A$ and that the vector of authority scores is the Perron eigenvector of AA^T . The construction of HITS's focused subgraph is a combination of text content relevancy with the query and of hyperlink considerations. Maximizing the probability of appearance of a web page on this subgraph is thus a composite problem out of the range of this thesis. We shall however study the optimization of HITS authority, for a given focused subgraph.

There exist several variants of HITS. First of all, as the matrix $A^T A$ is not necessarily irreducible, one can add a small positive number to this matrix [LM05a] in order to force irreducibility. Then the HITS vector is uniquely defined and positive. Another possibility is to consider the matrix exponential of the adjacency matrix and rank web pages according to the Perron vector of $\exp(A)^T \exp(A)$. This Exponentiated HITS [MRS⁺01] also guarantees the uniqueness of the Perron vector. Another variant of HITS is the Center rank presented in [BGH⁺04]. The goal of the Center rank is to define a score in the middle between hub and authority scores. In practice, the Center rank is given by the Perron eigenvector of the matrix $A^T A + AA^T$.

The SALSA algorithm [LM00] shares the same first step as HITS. Then we normalize the rows of A^T and A to get matrices A_c^T and A_r . The SALSA authority score is defined as the invariant measure of the stochastic matrix $A_c^T A_r$. In fact, with natural assumptions, this measure is proportional to the indegree of the web page. The authors show that the interest of the ranking algorithm lies in the combination of the two steps and not in one or the other alone. Thus from a hyperlink point of view, optimizing the rank in SALSA simply consists in maximizing the number of hyperlinks pointing to the target page. We shall not study SALSA optimization any further.

PageRank, HITS and SALSA share a mutual reinforcement property [DHH⁺02]: good hubs are pages that link to good authorities and good authorities are linked to by good hubs. For PageRank, the hub and authority scores are the same and interpreted as the popularity of

the page. The relation between the hub and authority scores is given by a possibly normalized adjacency matrix. By mixing the features of those algorithms, it is possible to define new rankings, like Auth-Rank, Hub-Rank or Sym-Rank [DHH⁺02].

We also studied the optimization of Tomlin’s HOTS scores [Tom03]. In this case, the ranking is the vector of dual variables of an optimal flow problem. The flow represents an optimal distribution of web surfers on the web graph in the sense of entropy minimization. The dual variable, one by page, is interpreted as the “temperature” of the page, the hotter a page the better. The HOTS problem is indeed a modification of the matrix balancing problem, studied among others in [Har71, EHRS85, SZ90, Sch90], in order to address disconnected web graphs. Given a $n \times n$ nonnegative matrix A , the matrix balancing problem consists in finding a matrix X of the form $X = D^{-1}AD$ with D diagonal definite positive and such that $\sum_k X_{i,k} = \sum_j X_{j,i}$ for all i . Tomlin showed that the HOTS vector is solution of a nonlinear fixed point equation. It may be thus seen as a nonlinear eigenvector. Indeed, we show that most of the arguments available in the case of Perron vector optimization can be adapted to HOTS optimization.

The Sinkhorn ranking [Smi05, Kni08] is based on the equivalence scaling problem. Given an $m \times n$ nonnegative matrix A , we search for a matrix X of the form $X = D^L A D^R$ with D^L and D^R diagonal definite positive and such that X is bistochastic. The Sinkhorn-Knopp [KS67] algorithm is a famous algorithm designed for the resolution of the scaling problem. It can be written as

$$D_{i,i}^L \leftarrow \frac{1}{\sum_{j \in [m]} A_{i,j} D_{j,j}^R}, \quad \forall i \in [m]$$

$$D_{j,j}^R \leftarrow \frac{1}{\sum_{i \in [n]} D_{i,i}^L A_{i,j}}, \quad \forall j \in [n]$$

Smith [Smi05] proposed to rank web pages according to $D_{i,i}^R / D_{i,i}^L$. This ranking satisfies the following reversal symmetry: the Sinkhorn ranking for A^T is the inverse of the Sinkhorn ranking for A . Knight [Kni08] proposed an alternative interpretation. He remarks that scaling the adjacency matrix so that it is doubly stochastic is equivalent to scaling it so that its stationary distribution is uniform. Indeed, $D^L A D^R$ is stochastic, so it has a stationary distribution, which can be then interpreted as in the PageRank model as a web surfer distribution. Then if Page i has a small $D_{i,i}^R$ value, this means that it tends to emit traffic, while if Page i has a small $D_{i,i}^L$ value, this means that it tends to draw traffic. He thus considers $1/D_{i,i}^R$ as the hub score of Page i and $1/D_{i,i}^L$ as the authority score of Page i . This ranking has also been proposed in [GLM09] starting from an offense-defense model.

Akian, Gaubert and Ninove [AGN06] gave a self-validating web ranking called T-PageRank. Then the random surfer does not click on the hyperlinks on the current page with a uniform probability but he tends to follow more links that lead to pages with a higher score. A temperature parameter controls the confidence web surfers put on the web ranking, the higher the temperature, the less confident web surfers are. The authors show that for sufficiently big temperatures, the T-PageRank is uniquely defined and can be computed by an iterative fixed point iterative scheme. They also show that if web surfers are too confident in the web ranking, then the T-PageRank is not unique and the limit of the fixed point scheme actually depends on the initial condition.

In [DL07], Delvenne and Libert gave thermodynamical arguments to rank web pages according to the elementwise product of the left and right Perron eigenvectors of the adjacency

matrix, or of an irreducible modification of the adjacency matrix.

2.4 Nonlinear Perron-Frobenius theory

The common point between most of these ranking algorithms is the so-called (nonlinear) Perron-Frobenius theory. This theory deals with monotone and homogeneous maps. It has a multiplicative and an additive formulation

Definition 2.1. A map $T : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ is monotone if for all vectors p, q such that $p \leq q$, $T(p) \leq T(q)$. A map $T : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ is homogeneous if for all vector p and for all nonnegative real λ , $T(\lambda p) = \lambda T(p)$.

Definition 2.2. A map $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is additively homogeneous if for all vector p and for all real λ , $T(\lambda + p) = \lambda + T(p)$.

We can transform a multiplicative monotone, homogeneous map T^\times into a monotone, additively homogeneous map T^+ and vice versa by the following operation called the “logarithmic glasses”:

$$T^+(p) = \log(T^\times(\exp(p)))$$

where \log and \exp act elementwise.

The following results show that monotone and nonexpansive maps are indeed nonexpansive. Hence, they are well suited for iterative algorithms.

Proposition 2.4 ([CT80]). *An additively homogeneous map is nonexpansive for the sup-norm if and only if it is monotone.*

For a more general result, we shall need Hilbert’s projective metric.

Definition 2.3. For x, y two vectors of \mathbb{R}^n , Hilbert’s projective metric between x and y is defined as

$$d(x, y) = \log\left(\max_{i,j \in [n]} \frac{x_i y_j}{y_i x_j}\right)$$

Proposition 2.5 ([Bus73]). *Any monotone and homogeneous map is nonexpansive in Hilbert’s metric.*

Definition 2.4. For a map $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ or $T : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$, we call the graph of T and we denote it $G(T)$, the directed graph with nodes $1, \dots, n$ and an arc from i to j if and only if $\lim_{t \rightarrow +\infty} T(te_i) = +\infty$ where e_i is the i th basis vector.

The following results give conditions for the existence and uniqueness of the “eigenvector” of a monotone, (additively or multiplicatively) homogeneous map.

Theorem 2.1 (Theorem 2 in [GG04b]). *Let T be a monotone, additively homogeneous map. If $G(T)$ is strongly connected, then there exists $u \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$ such that $T(u) = \lambda + u$. We say that u is an additive eigenvector of f .*

The ranking p then is defined as a (nonlinear) Perron eigenvector

$$T(p) = \mu p$$

where the operator $T : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ is monotone and homogeneous. For instance, T can be the linear operator corresponding to an entrywise nonnegative matrix.

Theorem 2.2 (Corollary 2.5 in [Nus88], Theorem 2.3 in [FGH11], Theorem 6.8 in [AGN12]). *Let $T : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ be a continuously differentiable map which is monotone and homogeneous and has an eigenvector p with only positive entries. If $\nabla T(p)$ is irreducible, then there is at most one eigenvector in $(\mathbb{R}_+ \setminus \{0\})^n$ up to a multiplicative factor. If $\nabla T(p)$ is primitive, then all the orbits defined by*

$$p_{k+1} = \frac{T(p_k)}{\|T(p_k)\|}$$

for a given norm $\|\cdot\|$ converge to $\frac{p}{\|p\|}$ linearly at a rate equal to $\frac{|\lambda_2|}{\rho}$, the ratio of the second and largest eigenvalues of ∇T .

We can also write this theorem in additive form.

Theorem 2.3 (Corollary 2.5 in [Nus88], Theorem 2.3 in [FGH11], Theorem 6.8 in [AGN12]). *Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a continuously differentiable map which is monotone and additively homogeneous and has an additive eigenvector $p \in \mathbb{R}^n$. If $\nabla T(p)$ is irreducible, then the eigenvector is unique up to an additive factor. If $\nabla T(p)$ is primitive, then all the orbits defined by*

$$p_{k+1} = T(p_k) - \psi(T(p_k))$$

for a given additively homogeneous function $\psi : \mathbb{R}^n \rightarrow \mathbb{R}$ converge to $p - \psi(p)$ linearly at a rate equal to $|\lambda_2(\nabla T(p))| = \max\{|\lambda|; \lambda \in \text{spectrum}(\nabla T(p)), \lambda \neq 1\}$.

These theorems have been stated with more general assumptions, among others semi-differentiability [AGN12] and infinite state space [Nus88]. However, for the sake of simplicity, we present them here in this simpler form.

These results give a general framework to prove that the rankings are well defined, i.e. that the score is positive and unique, and that the power algorithm used to compute them indeed converges to the expected ranking.

While PageRank optimization has been studied by several authors, the optimization of the scores obtained by the other web ranking algorithms is less well documented. This is the object of Chapter 7 below.

CHAPTER 3

Markov Decision Processes with implicit action spaces

Let us consider a web page such that a webmaster controls some of the hyperlinks. Then, an action consists in choosing a subset of the set of facultative hyperlinks. Unfortunately, there is then an exponential number of actions. Csáji, Jungers and Blondel [CJB10] have solved this problem by considering a graph rewriting that gives an equivalent problem with a polynomial number of states and actions. We propose in this section an alternative approach that does not involve a graph rewriting.

We assume that in each state, the action space can be described as a set that contains the extreme points of a polytope with a polynomial-time separation oracle. Then, the action sets are *implicitly* described and this description is concise in a computational sense. We call a Markov decision process with such a description of the action set and linear rewards and transitions a well-described Markov decision process. We prove that the infinite horizon average cost problem for well-described Markov decision processes is solvable in polynomial time (Theorem 3.1). This approach has the advantage that it does not involve any modification of the Markov decision process. Hence the proofs are simplified and the properties of the Markov decision process remain. For instance for the PageRank optimization problem, we shall see in Chapter 4 that we keep the uniform contraction factor available when computing the PageRank while we lose it with the graph rewriting approach.

In Section 3.1, we recall the main complexity results for finite Markov decision processes. In Section 3.2, we give the definition of a well-described Markov decision process and Theo-

rem 3.1. We also generalize the polynomial time solvability result to Markov decision processes with finite state space and convex action spaces in Section 3.3. We give in Theorem 3.2 a convex program, the solution of which gives the expected average cost (equal to the cycle time of the daily operator) and a bias vector of the problem.

3.1 Finite Markov Decision Processes

A finite *Markov decision process* is a 4-tuple $(I, (A_i)_{i \in I}, p, r)$ where I is a finite set called the *state space*; for all $i \in I$, A_i is the finite set of *admissible actions* in state i ; the function $p : I \times \cup_{i \in I}(\{i\} \times A_i) \rightarrow \mathbb{R}_+$ is the *transition law*, so that $p(j|i, a)$ is the probability to go to state j from state i when action $a \in A_i$ is selected; and $r : \cup_{i \in I}(\{i\} \times A_i) \rightarrow \mathbb{R}$ is the *reward function*, so that $r(i, a)$ is the instantaneous reward when action a is selected in state i .

Let $X_t \in I$ denote the state of the system at the discrete time $t \geq 0$. A *deterministic control strategy* ν is a sequence of actions $(\nu_t)_{t \geq 0}$ such that for all $t \geq 0$, ν_t is a function of the history $h_t = (X_0, \nu_0, \dots, X_{t-1}, \nu_{t-1}, X_t)$ and $\nu_t \in A_{X_t}$. Of course, $\mathbb{P}(X_{t+1} = j | X_t, \nu_t) = p(j | X_t, \nu_t), \forall j \in [n], \forall t \geq 0$. More generally, we may consider *randomized* strategies ν where ν_t is a probability measure on A_{X_t} . A strategy ν is *stationary* (feedback) if there exists a function $\bar{\nu}$ such that for all $t \geq 0$, $\nu_t(h_t) = \bar{\nu}(X_t)$.

We may consider several optimization problems based on the framework of Markov decision processes. We consider an initial distribution μ representing the law of X_0 .

Given an integer T , the finite horizon Markov decision problem consists in maximizing

$$\mathbb{E}\left(\sum_{t=0}^{T-1} r(X_t, \nu_t)\right) \quad (3.1)$$

The total cost infinite horizon Markov decision problem consists in maximizing

$$\liminf_{T \rightarrow +\infty} \mathbb{E}\left(\sum_{t=0}^{T-1} r(X_t, \nu_t)\right) \quad (3.2)$$

Given $\alpha \in (0, 1)$, the discounted infinite horizon Markov decision problem consists in maximizing

$$\lim_{T \rightarrow +\infty} \mathbb{E}\left(\sum_{t=0}^{T-1} \alpha^t r(X_t, \nu_t)\right) \quad (3.3)$$

The average cost infinite horizon Markov decision problem, also called *ergodic control problem*, consists in maximizing

$$\liminf_{T \rightarrow +\infty} \frac{1}{T} \mathbb{E}\left(\sum_{t=0}^{T-1} r(X_t, \nu_t)\right) \quad (3.4)$$

In all four cases, the maximum is taken over the set of randomized control strategies ν . In fact, for the total cost, discounted and average cost problems, the supremum attained by randomized (or even deterministic) stationary feedback strategies (Theorem 9.1.8 in [Put94] for instance).

A Markov decision process is *unichain* if the transition matrix corresponding to every stationary policy has a single recurrent class. Otherwise it is *multichain*. When the problem is unichain, the value of the average cost problem does not depend on the initial distribution

whereas when it is not, one may consider a vector $(g_i)_{i \in I}$ where g_i represents the value of the problem (3.4) when starting from state i . The total cost infinite horizon problem might not have a finite value but it does have a finite value if the mean expected reward in each recurrent class is 0.

We now state a classical complexity result for Markov decision problems

Proposition 3.1 (Theorem 1 in [PT87]). *The finite horizon, discounted and average cost Markov Decision Process problems are P-complete, and hence polynomial time solvable.*

For the finite horizon problem, the dynamic programming algorithm finds the value of the problem and optimal strategies in a time proportional to nT . For the discounted problem, the value iterations algorithm [Tse90] and the policy iterations algorithm [Put94] finish in polynomial time. Strongly polynomial time algorithms also exist for the discounted problem [Ye05, Ye11]. For the average cost problem, one may consider a linear programming reformulation of the Markov decision problem (see Theorem 9.3.8 in [Put94] for instance) in order to prove its polynomial time solvability. It is also a simple corollary of the polynomial time solvability of the average cost problem that the total cost infinite horizon problem is solvable in polynomial time.

The policy iterations algorithm has received much attention these years following [Fea10], that gave an exponential lower bound for the policy iterations on the average cost Markov decision problem. This result has motivated the paper [HDJ11] that raises the question whether policy iterations for PageRank optimization is a polynomial time algorithm or not.

3.2 Polynomial time solvability of well-described Markov decision problems

We prove in this section, corresponding to Section IV-A in [FABG13], a general result of independent interest concerning Markov decision processes with implicitly defined action sets. We introduce the notion of *well-described* Markov decision processes, in which, although there may be an *exponential* number of actions, there is a polynomial time *strong separation oracle* for the actions polytope.

In fact, classical complexity results, like Theorem 1 in [PT87], assume that the actions are explicitly enumerated [PT87]. As polynomial time means, as usual, polynomial in the input length (number of bits of the input) and the input includes the description of the actions sets, this result only leads to an exponential bound in the PageRank optimization case, where there is an exponential number of actions.

We prove in Theorem 3.1 below that the polynomial time solvability of ergodic control problems subsists when the action sets are implicitly defined, even in the multi-chain framework. This is based on the combinatorial developments of the theory of Khachiyan’s ellipsoid method, by Groetschel, Lovász and Schrijver [GLS88]. We refer the reader to the latter monograph for more background on the notions of strong separation oracles and well described polyhedra. We note that maximization or separation oracles have been previously considered in dynamic programming for different purposes (dealing with unknown parameters [GLD00, XM10], or approximating large scale problems [KHG06]).

Definition 3.1 (Def. 6.2.2 of [GLS88]). We say that a polyhedron \mathcal{B} has *facet-complexity* at most ϕ if there exists a system of inequalities with rational coefficients that has solution set \mathcal{B}

and such that the encoding length of each inequality of the system (the sum of the number of bits of the rational numbers appearing as coefficients in this inequality) is at most ϕ .

A *well-described polyhedron* is a triple $(\mathcal{B}; n, \phi)$ where $\mathcal{B} \subseteq \mathbb{R}^n$ is a polyhedron with facet-complexity at most ϕ . The *encoding length* of \mathcal{B} is by definition $n + \phi$.

Definition 3.2 (Problem (2.1.4) of [GLS88]). A *strong separation oracle* for a set K is an algorithm that solves the following problem: given a vector y , decide whether $y \in K$ or not and if not, find a hyperplane that separates y from K ; i.e., find a vector c such that $c^T y > \max\{c^T x, x \in K\}$.

Inspired by Definition 3.1, we introduce the following notion.

Definition 3.3. A finite Markov decision process $(I, (A_i)_{i \in I}, p, r)$ is *well-described* if for every state $i \in I$, we have $A_i \subset \mathbb{R}^{L_i}$ for some $L_i \in \mathbb{N}$, if there exists $\phi \in \mathbb{N}$ such that the convex hull of every action set A_i is a well-described polyhedron $(\mathcal{B}_i; L_i, \phi)$ with a polynomial time strong separation oracle, and if the rewards and transition probabilities satisfy $r(i, a) = \sum_{l \in [L_i]} a_l R_i^l$ and $p(j|i, a) = \sum_{l \in [L_i]} a_l Q_{i,j}^l$, $\forall i, j \in I, \forall a \in A_i$, where R_i^l and $Q_{i,j}^l$ are given rational numbers, for $i, j \in I$ and $l \in [L_i]$.

The *encoding length* of a well-described Markov decision process is by definition the sum of the encoding lengths of the rational numbers $Q_{i,j}^l$ and R_i^l and of the well-described polyhedra \mathcal{B}_i .

The situation in which the action spaces are given as usual in extension (by listing the actions) corresponds to the case in which A_i is the set of extreme points of a simplex Σ_{L_i} . The interest of Definition 3.3 is that it applies to more general situations in which the actions are not listed, but given implicitly by a computer program deciding whether a given element of \mathbb{R}^{L_i} is an admissible action in state i (the separation oracle). An example of such a separation oracle stems from the description of the convex hull of the set of admissible transition probabilities for the PageRank optimization problem (Theorem 4.1 in Chapter 4 below): here, a potential (randomized) action is an element of \mathbb{R}^n , and to check whether it is admissible, it suffices to check whether one of the inequalities in (4.11) is not satisfied.

Theorem 3.1. *The average cost infinite horizon problem for a well-described (multichain) Markov decision process can be solved in a time polynomial in the input length.*

Proof. We shall use the notations of Definition 3.3. Consider the polyhedron \mathcal{G} consisting of the couples of vectors $(v, g) \in \mathbb{R}^I \times \mathbb{R}^I$ satisfying the constraints

$$\begin{aligned} g_i &\geq \sum_{j \in I} \sum_{l \in [L_i]} a_l Q_{i,j}^l g_j, \quad \forall i \in I, a \in A_i \\ v_i + g_i &\geq \sum_{l \in [L_i]} a_l R_i^l + \sum_{j \in I} \sum_{l \in [L_i]} a_l Q_{i,j}^l v_j, \quad \forall i \in I, a \in A_i. \end{aligned} \tag{3.5}$$

Theorem 9.3.8 in [Put94] implies that the average cost problem reduces to minimizing the linear form $(v, g) \mapsto \sum_{j \in I} g_j$ over \mathcal{G} . Every optimal solution (v, g) of this linear program is such that g_j is the optimal mean payment per time unit starting from state j . We recover optimal strategies of the ergodic problem through dual optimal solution of the linear program.

By Theorem 6.4.9 in [GLS88], we know that a linear program over a well-described polyhedron with a polynomial time strong separation oracle is polynomial time solvable. Moreover,

Theorem 6.5.14 in [GLS88] asserts that we can find a dual optimal solution in polynomial time.

Let us construct such an oracle for \mathcal{G} . Given a point $(g, v) \in \mathbb{Q}^n \times \mathbb{Q}^n$, compute for all $i \in I$: $\max_{a \in \text{co}(A_i)} \sum_{l \in [L_i]} a_l (\sum_{j \in I} Q_{i,j}^l g_j) - g_i$ and $\max_{a \in \text{co}(A_i)} \sum_{l \in [L_i]} a_l (R_i^l + \sum_{j \in I} Q_{i,j}^l v_j) - v_i - g_i$. Those problems are linear problems such that, by hypothesis, we have a polynomial time strong separation oracle for each of the well-described polyhedral admissible sets $\mathcal{B}_i = \text{co}(A_i)$. Thus they are polynomial time solvable. If the $2n$ linear programs return a nonpositive value, then this means that (g, v) is an admissible point of (3.5). Otherwise, the solution a of any of those linear programs that have a negative value yields a strict inequality $g_i < \sum_{j \in I} \sum_{l \in [L_i]} a_l Q_{i,j}^l g_j$ or $v_i + g_i < \sum_{l \in [L_i]} a_l R_i^l + \sum_{j \in I} \sum_{l \in [L_i]} a_l Q_{i,j}^l v_j$. In both cases, the corresponding inequality determines a separating hyperplane.

To conclude the proof, it remains to check that the facet complexity of the polyhedron \mathcal{G} is polynomially bounded in the encoding lengths of the polyhedra \mathcal{B}_i and the rationals R_i^l and $Q_{i,j}^l$. Since the a_l 's appear linearly in the constraints (3.5), these constraints hold for all $a \in A_i$ if and only if they hold for all $a \in \mathcal{B}_i$ or equivalently, for all extreme points of \mathcal{B}_i . The result follows from Lemma 6.2.4 in [GLS88], which states that the encoding length of any extreme point of a well-described polyhedron is polynomially bounded in the encoding of the polyhedron. \square

Remark 3.1. This argument also shows that the discounted problem is polynomial time solvable.

3.3 Convex Markov decision processes

In this section, we give a generalization of the polynomial-time solvability result of the previous section to the infinite horizon expected average cost problem where the daily operator is a given convex, monotone and additively homogeneous function. In Section 3.2 we considered Markov decision processes with finite state space, polyhedral action spaces and linear rewards. Now, we shall accept any action space and rewards. We just need the state space to be finite and the daily operator to be effectively computable. For instance, if the action spaces are convex sets and the rewards are convex functions, then an ϵ -approximation of the daily operator is computable in polynomial time (Theorem 5.3.1 in [BTN01]).

We give in Theorem 3.2 a convex program, the solution of which gives the expected average cost, equal to the cycle time of the daily operator, and a bias vector of the problem. The proof relies on the existence of the cycle time of any monotone homogeneous and convex map [GG04a, Vig09a] and on inequalities of convexity.

Lemma 3.1. *Let f be an additively homogeneous convex function with Fenchel transform f^* . Let v_α be the unique vector such that $v_\alpha = f(\alpha v_\alpha)$, $V_\alpha(x) := \sup_{\pi} (1-\alpha) \sum_{m=0}^{+\infty} -\alpha^m \mathbb{E}_{\pi}^x [f_{x^m}^*(q^m)]$ and $V_n(x) := \sup_{\pi} \frac{1}{n} \sum_{l=0}^{n-1} -\mathbb{E}_{\pi}^x [f_{x^l}^*(q^l)]$ where the supremum is taken on the set the admissible measures π such that $\mathbb{P}_{\pi}^x(q^m \in \text{dom}(f_{x^m}^*)) = 1$. Then $V_\alpha = (1-\alpha)v_\alpha$ and $V_n(x) = \frac{1}{n} f^n(x)$.*

Proof. This derives from the interpretation of a monotone, additively homogeneous convex function as the daily operator of a Markov decision process with finite state space and convex action spaces [AG03]:

$$f_i(x) = \sup_{q^i \in \text{dom}(f_i^*)} \langle q^i, x \rangle - f_i^*(q^i), \quad \forall i \in [n] \quad \square$$

The following development is a simple generalization of Proposition 4 and Corollary 5 in Lehrer and Sorin's Tauberian theorem for dynamic programming article [LS92].

Lemma 3.2. *For any monotone, additively homogeneous convex function f , we have that $\forall \epsilon > 0, \forall N, \exists \alpha_0, \forall \alpha \geq \alpha_0, \forall x \in S, \exists n \geq N$ with $V_n(x) \geq V_\alpha(x) - \epsilon$.*

Proof. For all $n \in \mathbb{N} \cup \{\infty\}$ and for all $(a_m)_m$ bounded from above, we have

$$(1 - \alpha) \sum_{m=0}^n \alpha^m a_m = (1 - \alpha)^2 \sum_{m=0}^{n-1} \alpha^m (m+1) \left(\sum_{l=0}^m \frac{a_l}{l+1} \right) + (1 - \alpha) \alpha^n (n+1) \left(\sum_{l=0}^n \frac{a_l}{n+1} \right) \quad (3.6)$$

Let $M = \sup_{n,x} V_n(x)$. M is finite because $\limsup_n V_n = \chi(f)$ the cycle time of f [GG04a]. Given $\epsilon > 0$ and N , let α_0 be such that $\forall \alpha \leq \alpha_0$,

$$(1 - \alpha)^2 \sum_{m=0}^{N-1} \alpha^m (m+1) M \leq \epsilon/2 .$$

Given an $\epsilon/2$ -optimal strategy π for V_α , we apply (3.6) with $a_m = -\mathbb{E}_\pi^x[f_{x^m}^*(q^m)]$ and $n = +\infty$:

$$V_\alpha(x) - \epsilon/2 \leq \epsilon/2 + (1 - \alpha)^2 \sum_{m=N}^{+\infty} \alpha^m (m+1) \left(\sum_{l=0}^m \frac{a_l}{l+1} \right) ,$$

and as $(\sum_{l=0}^m \frac{a_l}{l+1}) \leq V_m(x)$, for all m , we get that a convex combination of $(V_m(x))_{m \geq N}$ is greater than $V_\alpha(x) - \epsilon$. This implies that at least one is greater (we denote it n), which yields the result. \square

Corollary 3.1. *For any monotone, additively homogeneous convex function f ,*

$$\limsup_{n \rightarrow +\infty} \frac{1}{n} f^n(x) \geq \limsup_{\alpha \rightarrow 1} (1 - \alpha) v_\alpha$$

Proof. Simple consequence of Lemma 3.2. \square

Let us denote \hat{f} the recession function of f defined by $\hat{f}(x) = \lim_{\lambda \rightarrow \infty} \frac{1}{\lambda} f(\lambda x)$ and $\chi(f)$ the cycle time of f defined by $\chi(f) = \lim_{n \rightarrow \infty} V_n(x) = \lim_{n \rightarrow \infty} \frac{1}{n} f^n(x)$.

Proposition 3.2 ([GG04a], Corollary 5.5.2 in [Vig09a]). *For any convex, monotone, additively homogeneous function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the cycle time of f exists and $\hat{f}(\chi(f)) = \chi(f)$.*

Theorem 3.2. *For any convex, monotone, additively homogeneous function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and for any $\beta \in \mathbb{R}^n$ such that $\beta > 0$, the optimal value of the convex optimization problem*

$$\begin{aligned} & \inf_{g \in \mathbb{R}^n, v \in \mathbb{R}^n} \beta \cdot g \\ & g \geq \hat{f}(g) \\ & g + v \geq f(v) \end{aligned}$$

is equal to $\beta \cdot \chi(f)$ and $g = \chi(f) = \lim_{n \rightarrow \infty} \frac{1}{n} f^n(x)$ at every optimum.

Proof. Let us denote C the convex set defined by the inequalities $g \geq \hat{f}(g)$ and $g + v \geq f(v)$. First of all, for all convex monotone additively homogeneous function f , $\chi(f)$ exists (Proposition 3.2).

We first prove that

$$\forall (g, v) \in C, \quad g \geq \chi(f) \quad (3.7)$$

For all $k \in \mathbb{N}$ and $w \in \mathbb{R}^n$

$$\begin{aligned} kg + f(w) &\geq \hat{f}(kg) + f(w) = \lim_{\lambda \rightarrow \infty} \frac{1}{\lambda} f(\lambda kg) + f(w) = \lim_{\lambda \rightarrow \infty} \frac{1}{\lambda} f(\lambda kg) + \frac{\lambda - 1}{\lambda} f(w) \\ &\geq \lim_{\lambda \rightarrow \infty} f\left(kg + \frac{\lambda - 1}{\lambda} w\right) = f(kg + w) \end{aligned}$$

Hence,

$$\begin{aligned} f(v) &\leq g + v \\ f^2(v) &\leq f(g + v) \leq g + f(v) \leq 2g + v \\ f^3(v) &\leq f(2g + v) \leq 2g + f(v) \leq 3g + v \end{aligned}$$

and by induction, we get that for all n , $f^n(v) \leq ng + v$, which implies that $g \geq \chi(f)$.

Now fix $\epsilon > 0$. Its is easy to see that $\hat{f}(\chi(f) + \epsilon) = \chi(f) + \epsilon$. We shall now just find $v \in \mathbb{R}^n$ such that

$$v + \chi(f) + \epsilon \geq f(v) .$$

For $\alpha \in [0, 1)$, let v_α be the unique solution of the discounted problem $f(\alpha v_\alpha) = v_\alpha$. By Corollary 3.1,

$$\chi(f) = \limsup \frac{1}{n} f^n(x) \geq \limsup (1 - \alpha)v_\alpha$$

We have

$$v_\alpha = f(\alpha v_\alpha) = f(v_\alpha - (1 - \alpha)v_\alpha) \geq f(v_\alpha) - Q(1 - \alpha)v_\alpha$$

for all $Q \in \partial f(v_\alpha)$ (∂f is the subdifferential of f defined as the Cartesian product of ∂f_i for $i \in [n]$ [AG03]). As $\limsup (1 - \alpha)v_\alpha \leq \chi(f)$, there exists α_0 such that for all $\alpha \geq \alpha_0$, $(1 - \alpha)v_\alpha \leq \chi(f) + \epsilon$. Now we have

$$\hat{f}_i(v_\alpha) = \sup_{p^i \in \text{dom} f_i^*} \langle p^i, v_\alpha \rangle \geq \langle q^i, v_\alpha \rangle$$

for all $q^i \in \partial f_i(v_\alpha)$, since for all x , $\partial f_i(x) = \{q^i \in \text{dom} f_i^* | \langle q^i, x \rangle - f_i^*(q^i) = f(x)\} \subseteq \text{dom} f_i^*$. By multiplicative homogeneity and monotonicity of \hat{f} we get

$$\chi(f) + \epsilon = \hat{f}(\chi(f) + \epsilon) \geq \hat{f}((1 - \alpha)v_\alpha) \geq (1 - \alpha)Qv_\alpha \geq f(v_\alpha) - v_\alpha$$

where Q can be any element of $\partial f(v_\alpha)$.

This inequality means that $(\chi(f) + \epsilon, v_\alpha) \in C$. As we can choose any arbitrary ϵ , for all $i \in [n]$, $\inf\{g_i, (g, v) \in C\} \leq \chi_i(f)$. The reverse inequality was given by (3.7), so the result holds. \square

CHAPTER 4

Ergodic control and Polyhedral approaches to PageRank optimization

4.1 Introduction

The PageRank introduced by Brin and Page [BP98] is defined as the invariant measure of a walk made by a random surfer on the web graph. When reading a given page, the surfer either selects a link from the current page (with a uniform probability), and moves to the page pointed by that link, or interrupts his current search, and then moves to an arbitrary page, which is selected according to given “zapping” probabilities. The rank of a page is defined as its frequency of visit by the random surfer.

The interest of the PageRank algorithm is to give each page of the web a measure of its popularity. It is a link-based measure, meaning that it only takes into account the hyperlinks between web pages, and not their content. It is combined in practice with content-dependent measures, taking into account the relevance of the text of the page to the query of the user, in order to determine the order in which the answer pages will be shown by the search engine. This leads to a family of search methods the details of which may vary (and are often not publicly known). However, a general feature of these methods is that among the pages with a comparable relevance to a query, the ones with the highest PageRank will appear first.

The importance of optimizing the PageRank, specially for e-business purposes, has led to the development of a number of companies offering Search Engine Optimization services. We refer in particular the reader to [CLF09] for a discussion of the PageRank optimization meth-

ods which are used in practice. Understanding PageRank optimization is also useful to fight malicious behaviors like link spamming, which intend to increase artificially the PageRank of a web page [GGM05a], [BYCL05].

The optimization of PageRank has been studied by several authors (see Section 2.2 for more details). Avrachenkov and Litvak analyzed in [AL06] the case of a single controlled page and determined an optimal strategy. In [MV06], Mathieu and Viennot established several bounds indicating to what extent the rank of the pages of a (multi-page) website can be changed, and derived an optimal referencing strategy in a special unconstrained case: if the webmaster can fix arbitrarily the hyperlinks in a web site, then, it is optimal to delete every link pointing outside the web site. To avoid such degenerate strategies, De Kerchove, Ninove and Van Dooren [dKNvD08] studied the problem of maximizing the sum of the PageRank coordinates in a web site, provided that from each page, there is at least one path consisting of hyperlinks and leading to an external page. They gave a necessary structural condition satisfied by an optimal outlink strategy. In [Nin08], Ninove developed a heuristic based on these theoretical results, which was experimentally shown to be efficient. In [IT09], Ishii and Tempo investigated the sensitivity of the PageRank to fragile (i.e. erroneous or imperfectly known) web data, including fragile links (servers not responding, links to deleted pages, etc.). They gave bounds on the possible variation of PageRank and introduced an approximate PageRank optimization problem, which they showed to be equivalent to a linear program. In [CJB10], (see also [CJB09] for more details), Csáji, Jungers and Blondel thought of fragile links as controlled links and gave an algorithm to optimize in polynomial time the PageRank of a single page.

Most of this chapter corresponds to [FABG13]. We study here a more general PageRank optimization problem, in which a webmaster, controlling a set of pages (her web site), wishes to maximize a utility function depending on the PageRank or, more generally, on the associated *occupation measure* (frequencies of visit of every link, the latter are more informative). For instance, the webmaster might wish to maximize the number of clicks per time unit of a certain hyperlink bringing an income, or the rank of the most visible page of her site, or the sum of the ranks of the pages of this site, etc. We consider specifically two versions of the PageRank optimization problem.

We first study a *continuous* version of the problem in which the set of actions of the webmaster is the set of admissible transition probabilities of websurfers. This means that the webmaster, by choosing the importance of the hyperlinks of the pages she controls (size of font, color, position of the link within the page), determines a continuum of possible transition probabilities. Although this model has been already proposed by Nemirovsky and Avrachenkov [NA08], its optimization does not seem to have been considered previously. This continuous version includes rather realistic constraints: for instance, the webmaster may start from a “template” or “skeleton” (given by designers), and be allowed to modify this skeleton only to a limited extent. Moreover, we shall allow *coupling constraints* between different pages (for instance, the rank of one page may be required to be greater than the rank of another page, constraints involving the sum of the PageRanks of a subset of pages are also allowed, etc.).

Following [IT09, CJB10], we also study a *discrete* version of the problem, in which in each page, there are obligatory links, facultative links and forbidden links. Then, the decision consists in selecting the subset of facultative links which are actually included in the page.

We show that when there are no coupling constraints between different pages and when the utility function is linear, the continuous and discrete problems both can be solved in

polynomial time by reduction to a linear program (our first main result, Theorem 4.3). When specialized to the discrete problem, this extends Theorem 1 of [CJB10], which only applies to the case in which the utility function represents the PageRank of a single page. The proof of Theorem 4.3 relies on the observation that the polytope generated by the transition probability measures that are uniform on some subsets of pages has a concise representation with a polynomial number of facets (Theorem 4.1). Then, Theorem 4.3 follows as a direct corollary of our works on Markov decision processes with implicitly defined action spaces (Theorem 3.1).

Proposition 4.7 yields a fixed point scheme with a contraction rate independent of the number of pages. Indeed, the contraction rate depends only on the “damping factor” (probability that the user interrupts his current search). Therefore, this problem can be solved efficiently for very large instances by Markov decision techniques. Our results show that optimizing the PageRank is not much more difficult than computing it, provided there are no coupling constraints: indeed, Proposition 4.9 shows that by comparison, the execution time is only increased by a $\log n$ factor, where n is the number of pages. Note that the Markov decision process which we construct here is quite different from the one of [CJB10], the latter is a stochastic shortest path problem, whose construction is based on a graph rewriting technique, in which intermediate (dummy) nodes are added to the graph. Such nodes are not subject to damping and therefore, the power iteration loses its uniform contraction. In our approach, we use a more general ergodic control model, which allows us to consider a general linear utility function, and avoids adding such extra nodes. Experiments also show that the present approach leads to a faster algorithm (Section 4.7.2).

We also study the continuous problem with general (linear) coupling constraints, and show that the latter can also be solved in polynomial time by reduction to a *constrained* ergodic control problem. Proposition 4.15 yields an algorithm to solve the PageRank optimization problem with coupling constraints, which scales well if the number of coupling constraints remains small. The resolution uses Lagrangian relaxation and convex programming techniques like the bundle method. There is little hope to solve efficiently, in general, the discrete problem with general coupling constraints since Csáji, Jungers and Blondel have proved in [CJB10] that the discrete PageRank optimization problem with mutual exclusion constraints is NP-complete. Nevertheless, we develop a heuristic for the discrete PageRank optimization problem with linear coupling constraints, based on the optimal solution of a relaxed continuous problem (Section 4.7.3). On test instances, approximate optimality certificates show that the solution found by the heuristic is at most at 0.4% of the optimum.

Using the concept of mean reward before teleportation, we identify in Theorem 4.4 (our second main result) assumptions under which there exists a “master” page to which all controlled pages should point. The theorem gives an ordering of the pages such that in loose terms, the optimal strategy is at each page to point to the allowed pages with highest order. The structure of the obtained optimal website is somehow reminiscent of Theorem 12 in [dKNvD08], but in [dKNvD08], there is only one constraint: the result is thus different.

When the problem has coupling constraints, the mean reward before teleportation still gives information on optimal strategies (Theorem 4.5).

We report numerical results on the web site of one of my co-authors [FABG13] (including an aggregation of surrounding pages) as well as on a fragment of the web (4.10^5 pages from the universities of New Zealand).

We finally note that an early Markov Decision Model for PageRank optimization was introduced by Bouhtou and Gaubert in 2007, in the course of the supervision of the student

project of Vlasceanu and Winkler [WV07].

The chapter is organized as follows. In Section 4.2, we introduce the general PageRank optimization problem. In Section 4.3.2, we show how the PageRank optimization problems reduce to ergodic control and we give a concise description of the polytope of uniform transition probabilities. In Section 4.4, we show that the continuous and discrete PageRank optimization problems with local constraints are solvable in polynomial time (Theorem 4.3). Section 4.4.2 describes an efficient fixed point scheme for the resolution of the PageRank optimization problem with local constraints. In Section 4.5, we give the “master page” Theorem (Theorem 4.4). We deal with coupling constraints in Section 4.6. We give experimental results on real data in Section 4.7.

4.2 PageRank optimization problems

4.2.1 Optimization of PageRank

The problem we are interested in is the optimization of PageRank. We study two versions of this problem. In the continuous PageRank Optimization problem, the webmaster can choose the importance of the hyperlinks of the pages she controls and thus she has a continuum of admissible transition probabilities (determined for instance by selecting the color of a hyperlink, the size of a font, or the position of a hyperlink in a page). This continuous model is specially useful in e-business applications, in which the income depends on the effective frequency of visit of pages by the users, rather than on its approximation provided by Google’s PageRank. The Continuous PageRank Optimization Problem is given by:

$$\max_{\pi, P} \{U(\pi, P) ; \pi = \pi P, \quad \pi \in \Sigma_n, \quad P \in \mathcal{P}\} \quad (4.1)$$

Here, $\Sigma_n := \{x \in \mathbb{R}^n \mid x_i \geq 0, \forall i \in [n]; \sum_{i \in [n]} x_i = 1\}$ is the *simplex* of dimension n , U is a *utility* function and \mathcal{P} is a set representing the set of all admissible transition probability matrices. We denote by P_i the i th row of a matrix P . We shall distinguish *local constraints*, which can be expressed as $P_i \in \mathcal{P}_i$, where $\mathcal{P}_i \subset \Sigma_n$, is a given subset, and *global constraints*, which couple several vectors $P_{i..}$. Thus, local constraints only involve the outlinks from a single page, whereas global constraints involve the outlinks from different pages. Then, $\mathcal{P} \subseteq \prod_{i \in [n]} \mathcal{P}_i$ with equality when there are only local constraints. We shall consider the situation in which each \mathcal{P}_i is a polytope (or more generally an effective convex set) with for all i , following (2.2),

$$\mathcal{P}_i = \alpha \mathcal{S}_i + (1 - \alpha)z, \quad \mathcal{S}_i \subset \Sigma_n. \quad (4.2)$$

If we restrict our attention to Google’s PageRank (with uniform transition probabilities), we arrive at the following combinatorial optimization problem. For each page i , as in [IT09] and [CJB10], we partition the set of potential links (i, j) into three subsets, consisting respectively of *obligatory links* \mathcal{O}_i , *prohibited links* \mathcal{I}_i and the set of *facultative links* \mathcal{F}_i . Then, for each page i , we must select the subset J_i of the set of facultative links \mathcal{F}_i which are effectively included in this page. Once this choice is made for every page, we get a new webgraph, and define the transition matrix $S = S(J_1, \dots, J_n)$ as in(2.1). The matrix after teleportation is also defined as above by $P(J_1, \dots, J_n) := \alpha S(J_1, \dots, J_n) + (1 - \alpha)ez$. Then, the Discrete PageRank Optimization Problem is given by:

$$\max_{\pi, P} \{U(\pi, P) ; \pi = \pi P, \quad \pi \in \Sigma_n, \quad P = P(J_1, \dots, J_n), \quad J_i \subseteq \mathcal{F}_i, \quad i \in [n]\} \quad (4.3)$$

Remark 4.1. Problem (4.3) is a combinatorial optimization problem: if there are p_i facultative links in page i , the decision variable, (J_1, \dots, J_n) , takes 2^p values, where $p = p_1 + \dots + p_n$.

We shall be specially interested in the modeling of an income proportional to the frequency of clicks on some hyperlinks. Let $r_{i,j}$ be a reward per click for each hyperlink (i, j) . The latter utility can be represented by the following *linear utility function*, which gives the total income:

$$U(\pi, P) = \sum_{i \in [n]} \pi_i \sum_{j \in [n]} P_{i,j} r_{i,j} . \quad (4.4)$$

Unless stated otherwise, we will consider the total income linear utility in the sequel.

Remark 4.2. The problem of maximizing the total PageRank of a web site (sum of the PageRanks of its pages) is obtained as a special case of (4.4). Indeed, if this web site consists of the subset of pages $I \subseteq [n]$, one can set $r_{i,j} = \chi_I(i)$, $\forall i, j \in [n]$, where χ_I is the characteristic function of I (with value 1 if $i \in I$ and 0 otherwise). Then

$$U(\pi, P) = \sum_i \pi_i \sum_j P_{i,j} r_{i,j} = \sum_{i \in I} \pi_i .$$

Remark 4.3. Note that the general form of the utility function assumes that we receive the same instantaneous reward $r_{i,j}$ when the surfer follows the hyperlink (i, j) and when the surfer stops the current exploration at page i to teleport to page j . There is no loss of generality in assuming that it is so: assume that the surfer produces a reward of $r'_{i,j}$ when he follows the hyperlink (i, j) and 0 when he teleports to page j .

Using the fact that $\sum_{j \in [n]} r'_{i,j} z_j = \sum_{j \in [n]} \sum_{l \in [n]} r'_{i,l} z_l P_{l,j}$ and $P = \alpha S + (1 - \alpha) e z$, we show that $\alpha \sum_{i,j \in [n]} r'_{i,j} \pi_i S_{i,j} = \sum_{i,j \in [n]} (r'_{i,j} - (1 - \alpha) \sum_{l \in [n]} r'_{i,l} z_l) \pi_i P_{i,j}$. We then only need to set $r_{i,j} = r'_{i,j} - (1 - \alpha) \sum_{l \in [n]} r'_{i,l} z_l$.

We shall restrict our attention to situations in which π is uniquely defined for each admissible transition matrix $P \in \mathcal{P}$ (recall that this is the case in particular when $\alpha < 1$). Then the utility U is a function of P only.

Alternatively, it will be convenient to think of the utility as a function of the *occupation measure* $\rho = (\rho_{i,j})_{i,j \in [n]}$. The latter is the stationary distribution of the Markov chain (x_{t-1}, x_t) . Thus, $\rho_{i,j}$ gives the frequency of the move from page i to page j . The occupation measure ρ is a probability measure and it satisfies the flow relation, so that

$$\rho_{i,j} \geq 0, \quad \forall i, j \in [n], \quad \sum_{i,j \in [n]} \rho_{i,j} = 1, \quad \sum_{k \in [n]} \rho_{k,i} = \sum_{j \in [n]} \rho_{i,j}, \quad \forall i \in [n]. \quad (4.5)$$

The occupation measure may also be thought of as a matrix. Hence, we shall say that ρ is *irreducible* when the corresponding matrix is irreducible.

The occupation measure ρ can be obtained from the invariant measure π and the stochastic matrix P by $\rho_{i,j} = \pi_i P_{i,j}$, $\forall i, j \in [n]$ and, conversely, the invariant measure π can be recovered from ρ by $\pi_i = \sum_{j \in [n]} \rho_{i,j}$, $\forall i \in [n]$.

The map f which determines the stochastic matrix P from the occupation measure is given by:

$$P = f(\rho), \quad P_{i,j} = \frac{\rho_{i,j}}{\sum_k \rho_{i,k}}, \quad \forall i, j \in [n]. \quad (4.6)$$

Proposition 4.1. *The function f defined by (4.6) sets up a birational transformation between the set of irreducible occupation measures (irreducible matrices satisfying (4.5)) and the set of irreducible stochastic matrices. In particular, the Jacobian of f is invertible at any point of the set of irreducible occupation measures.*

Proof. As π is uniquely defined, its entries are a rational function of the entries of P (for instance, when P is irreducible, an explicit rational expression is given by Tutte's Matrix Tree Theorem [Tut01]). The invertibility of the Jacobian follows from the birational character of f . \square

This bijective correspondence will allow us to consider the occupation measure, rather than the stochastic matrix P , as the decision variable. Note that the utility function can be written as a linear function in terms of the occupation measure: $U(\pi, P) = \sum_{i,j \in [n]} \rho_{i,j} r_{i,j}$.

4.2.2 Design constraints of the webmaster

We now model the possible modifications made by the webmaster, who may be subject to constraints imposed by the designer of the web site (the optimization of the PageRank should respect the primary goal of the web site, which is in general to offer some content). We thus describe the set \mathcal{P} of admissible transition probabilities of (4.1).

Proposition 4.2. *Assume that $\mathcal{P} = \prod_{i \in [n]} \mathcal{P}_i$, that for all $i \in [n]$, \mathcal{P}_i is a closed convex set and that every matrix $P \in \mathcal{P}$ is irreducible. Then, the set \mathcal{R} of occupation measures arising from the elements of \mathcal{P} is also a closed convex set. Moreover, if every \mathcal{P}_i is a polytope, then so is \mathcal{R} .*

Proof. For all $i \in [n]$, \mathcal{P}_i is a closed convex set and so it is the intersection of a possibly infinite family of hyperplanes $(H_i^{(l)})_{l \in L}$. Every element P of $\prod_{i \in [n]} \mathcal{P}_i$ must satisfy the following inequalities, one for each $H_i^{(l)}$:

$$\sum_{j \in [n]} a_{i,j}^{(l)} P_{i,j} \leq b_i^{(l)}, \quad \forall i \in [n], \forall l \in L \quad (4.7)$$

Formulating these equalities in terms of the occupation measure ρ thanks to $P_{i,j} = \frac{\rho_{i,j}}{\sum_{j' \in [n]} \rho_{i,j'}}$ and Proposition 4.1, and rewriting Inequalities (4.7) in the form

$$\sum_{j \in [n]} a_{i,j}^{(l)} \rho_{i,j} \leq b_i^{(l)} \sum_{k \in [n]} \rho_{i,k}, \quad \forall i \in [n], \forall l \in L \quad (4.8)$$

we see that ρ satisfies a family of constraints of the form (4.8), together with the inequalities (4.5). Thus, \mathcal{R} is defined as the intersection of half-spaces and so, it is closed and convex.

The same argument shows that if for all $i \in [n]$, \mathcal{P}_i is a polytope, so is \mathcal{R} . \square

We next list some concrete examples of such inequalities.

Skeleton constraints Imagine that a designer gave a *skeleton* or *template* for page i . The latter may include a collection of mandatory sites to be pointed by page i . We shall abstract the skeleton by representing it by a fixed probability vector $q \in \Sigma_n$, giving the transition probabilities if no further hyperlinks are added. Assume now that the webmaster is allowed to modify the page for optimization purposes, as long as the hyperlinks she adds do not overtake the initial content of the web site. This can be modeled by requiring that no hyperlink included in the skeleton loses a proportion of its weight greater than μ . Such constraints can be written as $P_{i,j} \geq \alpha(1 - \mu)q_j + (1 - \alpha)z_j, \forall j \in [n]$.

Linear coupling constraints Constraints like the presence of specific outlinks *somewhere* on the pages of the website are non-local. Such constraints cannot be written simply in terms of the stochastic matrix P (because adding conditional probabilities relative to different pages makes little sense) but they can be written linearly in terms of the occupation measure ρ , $\sum_{i,j \in [n]} a_{i,j} \rho_{i,j} \leq b$, where the coefficients $a_{i,j}$ and b are given.

These constraints include for instance *coupling conditional probability constraints*, which can be written as: $\sum_{i \in I, j \in J} \rho_{i,j} \geq b \sum_{i \in I, k \in [n]} \rho_{i,k}$. This means that the probability for the random surfer to move to set J , given that he is now in set I , should not be smaller than b .

We may also need *effective frequency constraints*: the webmaster (ruling a set of pages I) may sign a contract with another website (ruling a set of pages J), promising to redirect to this site an effective proportion of the web traffic. Such constraints may be written as $\sum_{i \in I, j \in J} \rho_{i,j} \geq b$. However, we warn the reader that such a bold contract may lead to an unfeasible problem, unless b is properly chosen.

Combinatorial constraints In the discrete problem, one may wish to set combinatorial constraints like demanding the existence of a path between two pages or sets of pages [dKNvD08], setting mutual exclusion between two hyperlinks [CJB10] or limiting the number of hyperlinks [CJB10]. Such constraints may lead to harder combinatorial problems, the solution of which is however made easier by the polynomial-time solvability of a relaxed continuous problem (Section 4.7.3).

4.3 Reduction of the PageRank Optimization Problem with local constraints to Ergodic Control

We next show that the continuous and discrete versions of the PageRank optimization reduce to ergodic control problems in which the action sets are defined as extreme points of concisely described polyhedra. We shall see in Section 4.4 that such problems remain tractable even if the size of the action sets may be exponential.

4.3.1 Reduction of the Continuous PageRank Optimization Problem to Ergodic Control

For a polytope \mathcal{P} , we shall denote by $\text{extr}(\mathcal{P})$ the set of extreme points of \mathcal{P} .

Proposition 4.3. *Assume that there are only local constraints, i.e. $\mathcal{P} = \prod_{i \in [n]} \mathcal{P}_i$, that for all $i \in [n]$, \mathcal{P}_i is a polytope of the form (4.2) and that the utility function is an income proportional to the frequency of clicks (4.4). Then the continuous PageRank Optimization problem (4.1) is*

equivalent to the unichain ergodic control problem with finite state $[n]$, finite action set $\text{extr}(\mathcal{P}_i)$ in every state i , transition probabilities $p(j|i, a) = a_j$ and rewards $r(i, a) = \sum_{j \in [n]} r_{i,j} a_j$.

Proof. As $\alpha < 1$, $a \in \mathcal{P}_i$ implies $a_k > 0$ for all k . Thus the problem defined in the proposition is unichain. Randomized stationary strategies are of the form $\nu_t = \bar{\nu}(X_t)$ for some function $\bar{\nu}$ sending $i \in [n]$ to some element of $\mathcal{P}_i = \text{co}(\text{extr}(\mathcal{P}_i))$. To such a strategy is associated a transition matrix P of the websurfer, obtained by taking $P_{i,\cdot} = \bar{\nu}(i)$ and vice versa. Thus, the admissible transition matrices of the websurfer are admissible stationary feedback strategies.

Moreover, the ergodic theorem for Markov chains shows that when such a strategy is applied,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} r(X_t, \nu_t) \right) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} \sum_{j \in [n]} r_{X_t, j} \bar{\nu}_j(X_t) \right) = \sum_{i, j \in [n]} \pi_i P_{i,j} r_{i,j}$$

and so, the objective function of the ergodic control problem is precisely the total income. \square

Proposition 4.4. *Under the hypotheses of Proposition 4.3, the dynamic programming equation*

$$w_i + \psi = \max_{\nu \in \mathcal{P}_i} \nu(r_{i,\cdot} + w) \quad , \quad \forall i \in [n] \quad (4.9)$$

has a solution $w \in \mathbb{R}^n$ and $\psi \in \mathbb{R}$. The constant ψ is unique and is the value of continuous PageRank Optimization problem (4.1). An optimal strategy is obtained by selecting for each state i a maximizing $\nu \in \mathcal{P}_i$ in (4.9). The function w is often called the bias or the potential.

Proof. Theorem 8.4.3 in [Put94] applied to the unichain ergodic control problem of Proposition 4.3 implies the result of the proposition but with \mathcal{P}_i replaced by $\text{extr}(\mathcal{P}_i)$. But as the expression which is maximized is affine, using \mathcal{P}_i or $\text{extr}(\mathcal{P}_i)$ yields the same solution. \square

4.3.2 The polytope of uniform transition measures

In this section, we show that the Discrete PageRank Optimization problem (4.3) is equivalent to a relaxed (continuous) PageRank Optimization problem (4.1) (Theorem 4.2). For this we show that the polytope of uniform transition measures, the vertices of which represent the action space of the Discrete PageRank Optimization problem (4.3), admits a concise representation (Theorem 4.1).

We consider a given page i and we study the set of admissible transition probabilities from page i . With uniform transitions, this is a discrete set that we denote \mathcal{D}_i . For clarity of the explanation, we will write x_j instead of $S_{i,j}$ and write the proofs in the case $\alpha = 1$. In order to get back to $\alpha < 1$, we use the relation $P_{i,j} = \alpha S_{i,j} + (1 - \alpha)z_j$ (see Remark 4.5 at the end of this section). Indeed, even if the polytope of admissible transition probabilities is well defined for any $\alpha \leq 1$, the PageRank optimization problem requires $\alpha < 1$ to be well defined, so that we will indeed use Remark 4.5.

We partition the set of links from page i as the set of obligatory links \mathcal{O}_i , the set of prohibited links \mathcal{I}_i and the set of facultative links \mathcal{F}_i . Then, depending on the presence of obligatory links,

$$\mathcal{D}_i = \{q \in \Sigma^n \mid \mathcal{O}_i \subseteq \text{supp}(q) \subseteq \mathcal{O}_i \cup \mathcal{F}_i, q \text{ uniform probability measure on its support}\} \quad (4.10)$$

or if $\mathcal{O}_i = \emptyset$, it is possible to have no link at all and then to teleport with probability vector Z (usually Z is equal to the teleportation vector z but it may be different):

$$\mathcal{D}_i = \{q \in \Sigma^n \mid \text{supp}(q) \subseteq \mathcal{F}_i, q \text{ uniform probability measure on its support}\} \cup \{Z\}.$$

We study the polytope $\text{co}(\mathcal{D}_i)$, the convex hull of the discrete set \mathcal{D}_i . Although it is defined as the convex hull of an exponential number of points, we show that it has a concise representation.

Theorem 4.1. *If page i has at least one obligatory link, then the convex hull of the admissible discrete transition probabilities from page i , $\text{co}(\mathcal{D}_i)$, is the projective transformation of a hypercube of dimension $|\mathcal{F}_i|$ and, for any choice of $j_0 \in \mathcal{O}_i$, it coincides with the polytope defined by the following set of inequalities:*

$$\forall j \in \mathcal{I}_i, \quad x_j = 0 \quad \forall j \in \mathcal{F}_i, \quad x_j \leq x_{j_0} \quad (4.11a)$$

$$\forall j \in \mathcal{O}_i \setminus \{j_0\}, \quad x_j = x_{j_0} \quad \forall j \in \mathcal{F}_i, \quad x_j \geq 0 \quad (4.11b)$$

$$\sum_{j \in [n]} x_j = 1 \quad (4.11c)$$

Proof. Let \mathcal{S}_i be the polytope defined by Inequalities (4.11).

($\mathcal{D}_i \subseteq \mathcal{S}_i$): Let q a probability vector in \mathcal{D}_i : q is a uniform probability measure on its support and $\mathcal{O}_i \subseteq \text{supp}(q) \subseteq \mathcal{O}_i \cup \mathcal{F}_i$. As for all j in \mathcal{F}_i , $q_j \leq \frac{1}{|\text{supp}(q)|} = q_{j_0}$, q verifies the equalities.

($\text{extr}(\mathcal{S}_i) \subseteq \mathcal{D}_i$): Let us consider an extreme point x of \mathcal{S}_i . Inequalities (4.11b) and (4.11a) cannot be saturated together at a given coordinate $j \in \mathcal{F}_i$ because, if it were the case, then we would have $x_{j_0} = 0$ and thus $x = 0$, which contradicts $\sum_{j \in [n]} x_j = 1$.

We have $1 + |\mathcal{I}_i| + |\mathcal{O}_i| - 1$ independent equalities so the polytope is of dimension $|\mathcal{F}_i|$. To be an extreme point, x must thus saturate $|\mathcal{F}_i|$ inequalities. At every j in \mathcal{F}_i , Inequalities (4.11b) and (4.11a) cannot be saturated simultaneously (see the previous paragraph), so the only way to saturate $|\mathcal{F}_i|$ inequalities is to saturate one of (4.11b) or (4.11a) at every j in \mathcal{F}_i . Finally, x can only take two distinct values, which are 0 and $x_{j_0} = \frac{1}{|\text{supp}(x)|}$: it is a uniform probability on its support.

We then show that \mathcal{S}_i is the projective transformation ([Zie01], Section 2.6 for more background) of the hypercube H defined by the following set of inequalities:

$$\{\forall j \in \mathcal{I}_i, X_j = 0; \forall j \in \mathcal{O}_i, X_j = 1; \forall j \in \mathcal{F}_i, 0 \leq X_j \leq 1\} .$$

As $\mathcal{O}_i \neq \emptyset$, H is embedded in the affine hyperplane $\{X \in \mathbb{R}^n \mid X_{j_0} = 1\}$. We can then construct the homogenization of H , $\text{homog}(H)$, which is the pointed cone with base H (see [Zie01] for more details). Finally \mathcal{S}_i is the cross-section of $\text{homog}(H)$ with $\{x \in \mathbb{R}^n \mid \sum_{j \in [n]} x_j = 1\}$. \square

The result of the theorem implies in particular that $\text{co}(\mathcal{D}_i)$ is combinatorially equivalent to a hypercube, i.e. that their face lattices are isomorphic [Zie01].

The next result concerns the case in which a page may have no outlink: it is necessary to consider this special case because then the websurfer teleports with probability Z_i to page i .




Figure 4.1: Projection of the polytope of uniform transition measures with one obligatory link ($|\mathcal{O}_i| = 1$) and three facultative links ($|\mathcal{F}_i| = 3$).

Proposition 4.5. *If page i has no obligatory link and if there exists $k \in \mathcal{I}_i$ such that $Z_k > 0$, then $\text{co}(\mathcal{D}_i)$ is a simplex of dimension $|\mathcal{F}_i|$ defined by the following set of inequalities:*

$$\sum_{j \in [n]} x_j = 1, \quad x_k \geq 0 \quad (4.12a)$$

$$\forall j \in \mathcal{I}_i \setminus \{k\}, \quad x_j = \frac{Z_j}{Z_k} x_k, \quad \forall j \in \mathcal{F}_i, \quad x_j \geq \frac{Z_j}{Z_k} x_k \quad (4.12b)$$

Proof. Let \mathcal{S}_i be the polytope defined by Inequalities (4.11).

($\mathcal{D}_i \subseteq \mathcal{S}_i$): Let q be an admissible discrete transition probability. Either q is a uniform probability measure on its support with $\text{supp}(q) \subseteq \mathcal{F}_i$ or $q = Z$. It is clear that Z satisfies all the inequalities. If $q \neq Z$, then $\forall j \in \mathcal{I}_i, q_j = 0$ (especially $q_k = 0$) and the equalities are straightforward to verify.

($\text{extr}(\mathcal{S}_i) \subseteq \mathcal{D}_i$): Now let x be an extreme point of the polytope defined by Inequalities (4.12). We have $1 + |\mathcal{I}_i| - 1$ independent equalities, so the polytope is of dimension $|\mathcal{F}_i|$. In order to be an extreme point, x must then saturate $|\mathcal{F}_i|$ inequalities. There are $|\mathcal{F}_i| + 1$ inequalities and by symmetry, we only need to consider two cases.

If Inequality (4.12a) is not saturated, all others have to be saturated. Then we have $x_j = \frac{Z_j}{Z_k} x_k, \forall j \in [n]$. As x sums to 1, $x = Z$.

If Inequality (4.12a) is saturated, then $x_k = 0$ and there exists $j \in \mathcal{F}_i$ such that for all $l \in \mathcal{F}_i \setminus \{j\}$ Inequality (4.12b) is saturated. As $x_k = 0$, this gives $\forall l \in \mathcal{F}_i \setminus \{j\}, x_l = 0$ and $x_j = 1$.

Finally, x is Z or x is an admissible discrete transition probability with $|\text{supp}(x)| = 1$. \square

Proposition 4.6. *If page i has no obligatory link and if for all $k \in \mathcal{I}_i, Z_k = 0$, then $\text{co}(\mathcal{D}_i)$ is the usual simplex of dimension $|\mathcal{F}_i| - 1$ with $x_j = 0, \forall j \in \mathcal{I}_i$.*

Proof. The extreme points of this simplex are clearly admissible discrete transition probabilities and the polytope contains every admissible discrete transition probabilities. \square

Remark 4.4. When there is no obligatory link, most of the admissible discrete transition probabilities are not extreme points of the polytope.

Remark 4.5. If we want to work with \mathcal{P}_i , the polytope of transition probabilities with damping factor α , we only need to use the relation $\mathcal{P}_i = \alpha \mathcal{S}_i + (1 - \alpha)z$ to get the actual inequalities. Hence, Theorem 4.1 applies also when $\alpha < 1$. For instance, $x_j = x_{j_0}$ remains but $x_j \geq 0$ becomes $x_j \geq (1 - \alpha)z_j$.

Theorem 4.2. *The discrete PageRank Optimization problem (4.3) is equivalent to a continuous PageRank Optimization problem (4.1) in which the action set \mathcal{P}_i is defined as in (4.2) by $\mathcal{P}_i = \alpha\mathcal{S}_i + (1 - \alpha)z$, ($0 < \alpha < 1$) and \mathcal{S}_i is one of the polytopes described in Theorem 4.1 or Proposition 4.5 or 4.6, depending on the presence of obligatory links.*

Proof. Arguing as in the proof of Proposition 4.3, we get that the discrete PageRank Optimization problem (4.3) is equivalent to an ergodic control problem with state space $[n]$, in which the action set in state i is the discrete set $\alpha\mathcal{D}_i + (1 - \alpha)z$ with \mathcal{D}_i defined in (4.10), and the rewards and transition probabilities are as in Proposition 4.3. The optimal solutions of the discrete PageRank Optimization problem coincide with the optimal stationary deterministic strategies. The analog of Equation (4.9) is now

$$w_i + \psi = \max_{\nu \in \text{aco}(\mathcal{D}_i) + (1 - \alpha)z} \nu(r_{i,\cdot} + w) \quad (4.13)$$

where $\text{co}(\mathcal{D}_i)$ is the convex hull of the set \mathcal{D}_i , i.e the polytope described in either Theorem 4.1 or Proposition 4.5 or 4.6. The polytope $\text{co}(\mathcal{D}_i)$ gives the transition laws in state i corresponding to randomized strategies in the former problem. Hence, the control problems in which the actions sets are $A_i = \alpha\mathcal{D}_i + (1 - \alpha)z$ or $\text{co}(A_i) = \alpha\text{co}(\mathcal{D}_i) + (1 - \alpha)z$ have the same value. Moreover, an optimal strategy of the problem with the latter set of actions can be found by solving (4.13) and selecting a maximizing action ν in (4.13). Such an action may always be chosen in the set of extreme points of $\text{co}(A_i)$ and these extreme points belong to A_i (beware however that some points of A_i may be not extreme). \square

4.4 Solving the PageRank Optimization Problem with local constraints

4.4.1 Polynomial time solvability of the PageRank Optimization Problem with local constraints

We have reduced the discrete and continuous PageRank Optimization problems to ergodic control problems in which the action sets are *implicitly defined* as the sets of extreme points of polytopes. Theorem 1 in [PT87] states that the ergodic control problem is solvable in polynomial time. However, in this result, the action sets are defined *explicitly*, whereas polynomial means, as usual, polynomial in the input length (number of bits of the input). Since the input includes the description of the actions sets, the input length is always larger than the sum of the cardinalities of the action sets. Hence, this result only leads to an exponential bound in our case (Remark 4.1).

However, by Theorem 3.1 in Chapter 3, we know that if we have a concise description of the action spaces, we can solve the average cost infinite horizon problem in polynomial time. Hence, as a consequence of Theorem 4.2, we get

Theorem 4.3. *If there are only local constraints, if the utility function is a rational total income utility (4.4) and if the teleportation vector and damping factor are rational, then the discrete problem (4.3) can be solved in polynomial time and the continuous problem (4.1) with well-described action sets (Definition 3.1) can also be solved in polynomial time.*

Proof. Thanks to Theorem 4.2, solving the continuous PageRank Optimization problem also solves the discrete PageRank Optimization problem. In addition, the coefficients appearing in the description of the facets of the polytopes of uniform transition measures are either 1, z_j or α and there are at most two terms by inequality (cf Section 4.3.2). This implies that these polytopes are well-described with an encoding length polynomial in the length of the input. Note also that we can find in polynomial time a vertex optimal solution of a linear program as soon as its feasible set is a well-described polytope as it is the case here (Lemma 6.5.1 in [GLS88]).

By Proposition 4.3, the ergodic control problem associated to a continuous PageRank Optimization problem with well-described action sets satisfies the conditions of Theorem 3.1 with $I = [n]$, $L_i = [n]$, $Q_{i,j}^l = \delta_{jl}$ and $R_i^l = r_{i,l}$ for $i, j \in [n], l \in L_i$. Thus it is polynomial time solvable. \square

Remark 4.6. In the previous theorem, the PageRank must be defined for any choice of facultative links. So we will assume that $\alpha < 1$. Nevertheless, even if $\alpha = 1$, we can still define the Markov Decision Process with finite state $[n]$, finite action set $\text{extr}(\mathcal{S}_i)$ in every state i , transition probabilities $p(j|i, a) = a_j$ and rewards $r(i, a) = \sum_{j \in [n]} r_{i,j} a_j$. This may be a multi-chain process, so that it may not have a unique invariant measure, but the solution of the infinite horizon average cost problem can be found in polynomial time by Theorem 3.1.

Theorem 3.1 and thus Theorem 4.3 is mostly of theoretical interest, since its proof is based on the ellipsoid algorithm, which is slow. We however give in Section 4.4.2 a fast scalable algorithm for the present problem.

Example 4.1. Consider again the graph from Figure 2.1, and let us optimize the sum of the PageRank scores of the pages of the site (colored). Assume that there are only local skeleton constraints (see Section 4.2.2): each page can change up to 20 % of the initial transition probabilities. The result is represented in Figure 4.2.

Example 4.2. We now consider a discrete PageRank optimization problem starting from the same graph. We set obligatory links to be the initial links and we represent them on the adjacency matrix in Figure 4.3 by squares. Facultative links are all other possible links from controlled pages.

4.4.2 Optimizing the PageRank via Value iteration

The PageRank optimization is likely not to be applied to the world wide web, but rather to a fragment of it, consisting of a web site (or of a collection of web sites of a community) and of related sites (see Section 4.5.1). However, even in such simplified instances, the number of design variables may be large, typically between thousands and millions. Hence, it is desirable to have scalable algorithms. We next describe two methods, showing that the optimization problem is computationally easy when there are no coupling constraints: then, optimizing the PageRank is essentially not more expensive than computing the PageRank.

Proposition 4.7. *Let T be the dynamic programming operator $\mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by*

$$T_i(w) = \max_{\nu \in \mathcal{S}_i} \alpha \nu(r_{i,\cdot} + w) + (1 - \alpha)z \cdot r_{i,\cdot}, \quad \forall i \in [n], \quad (4.14)$$

where $\mathcal{S}_i \subseteq \Sigma_n$ is closed. The map T is α -contracting and its unique fixed point w is such that $(w, (1 - \alpha)zw)$ is solution of the ergodic dynamic programming equation (4.9), with for all i , $\mathcal{P}_i = \alpha \mathcal{S}_i + (1 - \alpha)z$.

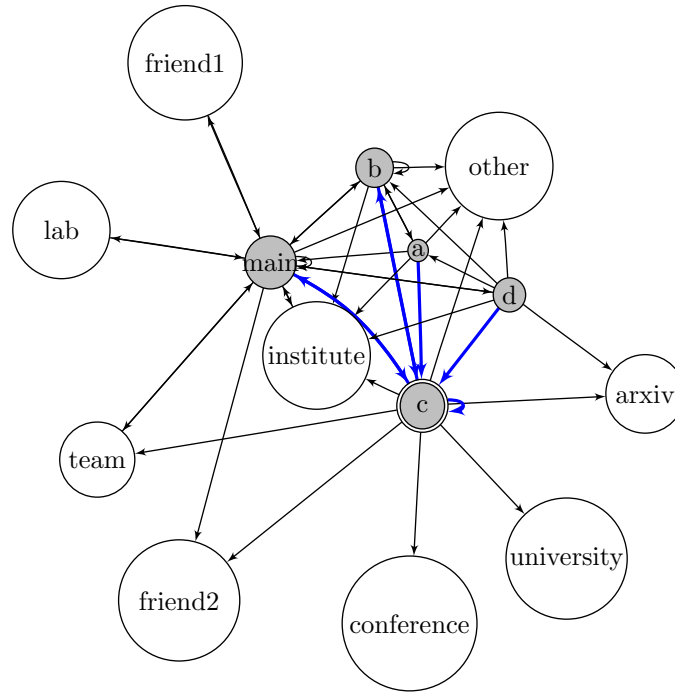


Figure 4.2: Web graph of Figure 2.1 optimized under local skeleton constraints. The optimal strategy consists in linking as much as possible to page "c" (actually, the page of a lecture), up to saturating the skeleton constraint. This page gains then a PageRank comparable to the one of the main page. The sum of the PageRank scores has been increased by 22.6%.




Figure 4.3: The web graph optimized under discrete uniform transitions constraints. In this case, the optimized graph has almost all internal links (links from a controlled page to another controlled page), so, for more readability, we display its adjacency matrix. The hyperlinks correspond to blue dots, obligatory links correspond to squares. The pages are ordered by decreasing average reward before teleportation (Section 4.5). The optimal strategy consists in adding a lot of internal links excluding certain pages, as will be explained by the master Page theorem below (Theorem 4.4).

Proof. The set \mathcal{S}_i is a set of probability measures so $\lambda \in \mathbb{R} \Rightarrow T(w + \lambda) = T(w) + \alpha\lambda$ and $v \geq w \Rightarrow T(v) \geq T(w)$. Thus T is α -contracting. Let w be its fixed point. For all $i \in [n]$,

$$w_i = \max_{\nu \in \mathcal{S}_i} \alpha \nu(r_{i,\cdot} + w) + (1 - \alpha)z(r_{i,\cdot} + w) - (1 - \alpha)zw = \max_{\nu' \in \mathcal{P}_i} \nu'(r_{i,\cdot} + w) - (1 - \alpha)zw$$

We get Equation (4.9) with constant $(1 - \alpha)zw$. \square

Remark 4.7. T is the dynamic programming operator of a total reward discounted problem with discount rate α and rewards $r'_{i,j} = r_{i,j} + \frac{1-\alpha}{\alpha} \sum_{l \in [n]} z_l r_{i,l}$ for transition from i to j (cf. Remark 4.3).

Remark 4.8. The fixed point found is just the mean reward before teleportation at the optimum (see Definition 4.1, Section 4.5).

We can then solve the dynamic programming equation (4.9) and so the PageRank Optimization Problem (4.1) or (4.3) with local constraints by value iteration (Algorithm 4.1). This algorithm consists in the fixed point iterations for the operator T . It converges to the fixed point w of T and $(w, (1 - \alpha)zw)$ is solution of the ergodic dynamic programming equation (4.9). The optimal linkage strategy is recovered by selecting the maximizing ν at each page. Thanks to the damping factor α , the iteration can be seen to be α -contracting. Thus the algorithm converges in a number of steps independent of the dimension of the web graph.

Algorithm 4.1 Computing the bias vector by value iteration

- 1: Start with an initial bias w^0 , choose a tolerance parameter ϵ , set $w^1 = T(w^0)$ and $k = 0$.
 - 2: **while** $\|w^k - w^{k+1}\| > \epsilon$ **do:** $w^{k+1} = T(w^k)$ (Eqn. (4.14)) and $k \leftarrow k + 1$
 - 3: **return** $w = w^k$, $\psi = (1 - \alpha)zw^k$
-

For the evaluation of the dynamic programming operator, one can use a linear program using to the description of the actions by facets. It is however usually possible to develop algorithms much faster than linear programming. We describe here a greedy algorithm for the discrete PageRank Optimization problem. The algorithm is straightforward if the set of obligatory links \mathcal{O}_i is empty (Propositions 4.5 and 4.6), so we only describe it in the other case. In Algorithm 4.2, J represents the set of facultative hyperlinks activated. We initialize it with the empty set and we augment it with the best hyperlink until it is not valuable any more to add a hyperlink.

Algorithm 4.2 Evaluation of the dynamic programming operator in the discrete problem

- 1: Initialization: $J \leftarrow \emptyset$ and $k \leftarrow 1$
 - 2: Sort $(w_l + r_{i,l})_{l \in \mathcal{F}_i}$ in decreasing order and let $\tau : \{1, \dots, |\mathcal{F}_i|\} \rightarrow \mathcal{F}_i$ be the sort function so that $w_{\tau(1)} + r_{i,\tau(1)} \geq \dots \geq w_{\tau(|\mathcal{F}_i|)} + r_{i,\tau(|\mathcal{F}_i|)}$.
 - 3: **while** $\frac{1}{|J|+|\mathcal{O}_i|} \sum_{l \in J \cup \mathcal{O}_i} (w_l + r_{i,l}) < w_{\tau(k)} + r_{i,\tau(k)}$ and $k \leq |\mathcal{F}_i|$ **do**
 - 4: $J \leftarrow J \cup \{\tau(k)\}$ and $k \leftarrow k + 1$
 - 5: **end while**
 - 6: $T_i(w) = \alpha \frac{1}{|J|+|\mathcal{O}_i|} \sum_{l \in J \cup \mathcal{O}_i} (w_l + r_{i,l}) + (1 - \alpha) \sum_{l \in [n]} z_l r_{i,l}$
-

Proposition 4.8. *When the constraints of the Discrete PageRank Optimization problem (4.3) are defined by obligatory, facultative and forbidden links, the greedy algorithm (Algorithm 4.2) started at page i returns $T_i(w)$ as defined in Proposition 4.7.*

Proof. The local constraints are obviously respected by construction. At the end of the loop, we have the best choice of facultative outlinks from page i with exactly $|J|$ outlinks. But as $\frac{1}{|J|+|\mathcal{O}_i|} \sum_{l \in J \cup \mathcal{O}_i} (w_l + r_{i,l}) \geq w_j + r_{i,j} \Leftrightarrow \frac{1}{|J|+|\mathcal{O}_i|} \sum_{l \in J \cup \mathcal{O}_i} (w_l + r_{i,l}) \geq \frac{1}{|J|+|\mathcal{O}_i|+1} \sum_{l \in J \cup \mathcal{O}_i \cup \{j\}} (w_l + r_{i,l})$, the sorting implies that we have the best choice of outlinks. \square

Remark 4.9. In order to handle upper and lower limits u_i and l_i on the number of links on Page i , we just need to replace $k \leq |\mathcal{F}_i|$ in Line 3 of Algorithm 4.2 by the condition $l_i \leq k + |\mathcal{O}_i| \leq \min(u_i, |\mathcal{O}_i| + |\mathcal{F}_i|)$.

Proposition 4.9. *An ϵ -approximation of the Discrete PageRank Optimization Problem (4.3) with only local constraints can be done in time*

$$O\left(\frac{\log(\epsilon)}{\log(\alpha)} \sum_{i \in [n]} |\mathcal{O}_i| + |\mathcal{F}_i| \log(|\mathcal{F}_i|)\right)$$

Proof. The value of the PageRank optimization problem is $(1 - \alpha)zw$ where $w = T(w)$. Thus it is bounded by $(1 - \alpha)\|z\|_1\|w\|_\infty = (1 - \alpha)\|w\|_\infty$. The greedy algorithm described in the preceding paragraph evaluates the i th coordinate of the dynamic programming operator T in a time bounded by $O(|\mathcal{O}_i| + |\mathcal{F}_i| \log(|\mathcal{F}_i|))$ (by performing a matrix-vector product and a sort). Thus it evaluates the dynamic programming operator in a time bounded by $O\left(\sum_{i \in [n]} |\mathcal{O}_i| + |\mathcal{F}_i| \log(|\mathcal{F}_i|)\right)$.

Now, if we normalize the rewards and if we begin the value iteration with $w^0 = 0$, the initial error is less than 1 in sup-norm. The fixed point iteration reduces this error by at least α , so we have to find $k \in \mathbb{N}$ such that $\alpha^k \leq \epsilon$. With $k \geq \frac{\log(\epsilon)}{\log(\alpha)}$, the result holds. \square

This should be compared to the complexity of PageRank computation by the power method [Ber05], which is $O\left(\frac{\log(\epsilon)}{\log(\alpha)} \sum_{i \in [n]} |\mathcal{O}_i| + |\mathcal{F}_i|\right)$.

4.4.3 A generalization of the algorithm of Csáji et al. for total income utility functions

We next describe an alternative method, which is based on an extension of an approach of Csáji, Jungers and Blondel in [CJB10]. The latter allows one to optimize the PageRank of a single page, by reduction to an ergodic Markov decision problem, after an augmentation of the state space. We will show that when the utility function of the PageRank to be optimized is linear, we can still solve it, and that value iterations is a possible algorithm.

The idea of [CJB10] is that if the control takes place on hyperlinks instead of transition probabilities, there are only two possible actions by facultative link: active or not.

We have the digraph $\mathcal{G} = ([n], \prod_{i \in [n]} \mathcal{O}_i)$ and a set of controlled links $\prod_{i \in [n]} \mathcal{F}_i$. Following [CJB10], we make an augmentation of the state space of the Markov chain. There are two cases for the state augmentation, depending on the presence of an obligatory hyperlink or not: see Figures 2.2 and 2.3 in Chapter 2.

We then add a cemetery state τ where web surfers go when they teleport. This is an analogy with the classical total reward point of view of discounted control problems. Remark that for every strategy, the associated Markov chain has one unique final class, the state t , where rewards are zero.

We denote V_1 for the set of states from the original model (squares on the figures), and V for the set of all states in the augmented Markov chain. We set nonzero instantaneous

rewards only for the links that lead to a state in V_1 . We build a big matrix R such that for all $j \in V_1$, $R_{i',j} = r'_{i,j}$ if $i' \in V_1$ is the image of $i \in V_1$ and $R_{k,j} = r'_{i,j}$ if node k represents the controlled hyperlink from page i to page j , where r' is defined as in Remark 4.7 by $r'_{i,j} = r_{i,j} + \frac{1-\alpha}{\alpha} \sum_{l \in [n]} z_l r_{i,l}$.

Proposition 4.10. *The discrete PageRank Optimization problem (4.3) where the utility function is an income proportional to the frequency of clicks on the hyperlinks (4.4) is equivalent to the total reward control problem for the controlled process $(Y_t)_{t \geq 0}$:*

$$\max_{\nu} \liminf_{T \rightarrow +\infty} \mathbb{E} \left(\sum_{t=0}^{T-1} R_{Y_t, Y_{t+1}} \right) \quad (4.15)$$

where

$$\mathbb{P}(Y_{t+1} = j | Y_t = i, \nu_t = a) = Q(i, j, a), \forall i, j \in V, \forall a \in A_i, \forall t \geq 0 ,$$

$$\mathbb{P}(Y_0 = j) = (1 - \alpha) z_j, \forall j \in V_1 ; \mathbb{P}(Y_0 = \tau) = \alpha ; \mathbb{P}(Y_0 = j) = 0, \forall j \notin V_1 \cup \{\tau\} ,$$

$Q(i, j, a)$ is the transition probability given by the state augmented model with action a chosen in the action set A_i and ν is a history dependent strategy.

Proof. Section 3.3 in [CJB10] gives the proof that the transition probabilities between original states are not altered by the state augmentation. For every strategy, the associated Markov chain has one unique final class, where where rewards are null: the total reward problem is then well defined. The rewards that we have defined are null except when we reach a state in V_1 and then they are equal in the two models.

The augmented state problem (4.15) admits optimal stationary strategies and it can be solved by the following dynamic programming equation (see [Put94] Theorem 10.4.3 for instance):

$$w_i = \max_{\nu \in A_i} \sum_{j \in V} Q(i, j, \nu) w_j + \sum_{j \in V} Q(i, j, \nu) R_{i,j}, \forall i \in V \quad \text{and} \quad w_\tau = 0$$

that has a unique solution w . We have simplified the second equation because here we only have one final class and the reward is 0 in it.

Because of the equivalence in transition probabilities and instantaneous rewards between initial and augmented state models, w_1 , the restriction of w to V_1 is solution of $T(w_1) = w_1$, where T is defined in Proposition 4.7 and Remark 4.7. The conclusion follows easily by Proposition 4.7. \square

The dynamic programming equation associated to Problem 4.15 has a number of states and actions in the order of $n + \sum_{i \in [n]} \mathcal{F}_i$. Evaluating its dynamic programming operator is even faster than with the algorithm of Proposition 4.9 because there is no need to sort. However, the value iteration derived from the augmented state model described here may need more iterations to converge because the discount factor does not act on all states: see Section 4.7.2 for a comparison on an example. We suggest it is better to solve this equation with policy iteration. See [Put94] for details of the algorithm in the total cost case and [HDJ11] for complexity issues of policy iterations on the graph augmented Markov decision process.

We have extended the results of [CJB10] for uniform transition probabilities to total income utility. Nevertheless, it may not be easy to extend it to the constrained case because Theorem 5.1 in [CJB10] shows that, if we accept mutual exclusion constraints, then the problem becomes NP-complete.

4.5 General shape of an optimized web site

We now use the previous model to identify the features of optimal link strategies. In particular, we shall identify circumstances under which there is always one “master” page, to which all other pages should link.

As in the work of De Kerchove, Ninove and Van Dooren [dKNvD08], we shall use the mean reward before teleportation to study the optimal outlink strategies.

Definition 4.1. Given a stochastic matrix P , the *mean reward before teleportation* is given by $v(P) := (I_n - \alpha S)^{-1} \bar{r}$, where $\bar{r}_i = \sum_j P_{i,j} r_{i,j}$.

Recall that S is the original matrix (without damping factor).

Proposition 4.11. *Suppose the instantaneous reward $r_{i,j}$ only depends on the current page i ($r_{i,j} = r'_i$). Denote $v(P)$ be the mean reward before teleportation (Definition 4.1). Then P is an optimal link strategy of the continuous PageRank Optimization problem (4.1) if and only if*

$$\forall i \in [n], \quad P_{i,\cdot} \in \arg \max_{\nu \in \mathcal{P}_i} \nu v(P)$$

Proof. We have $Pv(P) = v(P) - r' + \pi(P)r'$. Thus, using $\nu e = 1$, the condition of the proposition is equivalent to $\forall i \in [n], v_i(P) + \pi(P)r'_i = \max_{\nu \in \mathcal{P}_i} \nu(v(P) + r'_i e)$. By Proposition 4.4, this means that $v(P)$ is the bias of Equation (4.9) and that P is an optimal outlink strategy. \square

Remark 4.10. Proposition 4.11 shows that if P is any optimal outlink strategy, at every page i , the transition probability $P_{i,\cdot}$ must maximize the same linear function.

Remark 4.11. If two pages have the same constraint sets, then they have the same optimal outlinks, independently of their PageRank. This is no more the case with coupling constraints.

For the discrete PageRank Optimization problem, we have a more precise result:

Theorem 4.4 (Master Page). *Consider the Discrete PageRank Optimization problem (4.3) with constraints defined by given sets of obligatory, facultative and forbidden links. Suppose the instantaneous reward $r_{i,j}$ only depends on the current page i ($r_{i,j} = r'_i$). Let P be a transition matrix, $v = (I_n - \alpha S)^{-1} r$ is the mean reward before teleportation. Then P is an optimal link strategy if and only if for every controlled page i all the facultative links (i, j) such that $v_j > \frac{v_i - r'_i}{\alpha}$ are activated and all the facultative links (i, j) such that $v_j < \frac{v_i - r'_i}{\alpha}$ are deactivated. Any combination of facultative links such that $v_j = \frac{v_i - r'_i}{\alpha}$ can be activated.*

In particular, every controlled page should point to the page with the highest mean reward before teleportation, as soon as it is allowed to. We call it the “master page”.

Proof. Let P be an optimal strategy. By Remark 4.8, we know that the mean reward before teleportation at the optimum is a fixed point of the dynamic programming operator. In particular, it is invariant by the application of the greedy algorithm (Algorithm 4.2). Moreover, by Proposition 4.7, the mean reward before teleportation at the optimum is unique.

Thus, any optimal strategy must let the mean reward before teleportation invariant by the greedy algorithm. When there is no obligatory link from page i , either a link (i, j) is selected and $v_i = \alpha v_j + r'_i$ or no link is selected and $v_i = \sum_{k \in [n]} z_k v_k + r'_i > \alpha v_j + r'_i$ for all facultative link (i, j) . When there is at least one obligatory link, from Line 3 of the greedy algorithm, we know that, denoting J the set of activated links, all the links (i, j) verifying




Figure 4.4: Maximization of the sum of the PageRank values of the colored pages. Top: obligatory links; self links are forbidden; all other links from controlled pages are facultative. Bottom: bold arcs represent an optimal linking strategy. Page 4 points to all other controlled pages and Page 1, the master page, is pointed to by all other controlled pages. No facultative link towards an external page is selected.

$\frac{1}{|J|+|\mathcal{O}_i|} \sum_{l \in J \cup \mathcal{O}_i} v_l + r'_i < v_j + r'_j$, must be activated. This can be rewritten as $v_j > \frac{v_i - r'_i}{\alpha}$ because $v_i = \alpha \frac{1}{|J|+|\mathcal{O}_i|} \sum_{l \in J \cup \mathcal{O}_i} v_l + r'_i$.

Finally, activating any combination of the facultative links such that $v_j = \frac{v_i - r'_i}{\alpha}$ gives the same mean reward before teleportation.

Conversely, let P be an outlink strategy such that the mean reward before teleportation v verifies the hypotheses of the theorem. The equality $v = \alpha S v + r'$ can be rewritten as $(|J_i| + |\mathcal{O}_i|) \frac{v_i - r'_i}{\alpha} = \sum_{l \in J_i \cup \mathcal{O}_i} v_l$ for all $i \in [n]$. Launching the greedy algorithm on v shows that J_i is the best combination of outlinks with exactly $|J_i|$ outlinks and that adding or removing a link from it does not improve the value. Thus v is a fixed point of the dynamic programming operator and P is an optimal strategy. \square

The theorem is illustrated in Example 2 (Section 4.4) and Figure 4.4.

Example 4.3. The following simple counter examples show respectively that the conditions that instantaneous rewards only depend on the current page and that there are only local constraints are useful in the preceding theorem.

Take a two pages web graph without any design constraint. Set $\alpha = 0.85$, $z = (0.5, 0.5)$ and the reward per click $r = \begin{bmatrix} 1 & 10 \\ 2 & 2 \end{bmatrix}$. Then $v = (39.7, 35.8)$, Page 2 should link to Page 1 but Page 1 should link to Page 2 because $39.7 + 1 \leq 35.8 + 10$.

Take the same graph as in preceding example. Set $r' = (0, 1)$ and the coupling constraint that $\pi_1 \geq \pi_2$. Then every optimal strategy leads to $\pi_1 = \pi_2 = 0.5$. This means that there is no "master" page because both pages must be linked to in order to reach $\pi_i = 0.5$.

Remark 4.12. If every controlled page is allowed to point to every page, as in Figures 4.2 and 4.3, there is a master page to which every page should point. Actually, knowing that the optimal solutions are degenerate might be of interest to detect link spamming (or avoid being classified as a link spammer). The result of Proposition 4.11 and Theorem 4.4 can be related to [GGM05a], where the authors show various optimal strategies for link farms: patterns with every page linking to one single master page also appear in their study. We also remark that in [BYCL05], the authors show that making collusions is a good way to improve PageRank. We give here the page with which one should make a collusion.

Remark 4.13. If there exists a page with maximal reward in which all the hyperlinks can be changed, then this page is the master page. It will have a single hyperlink, pointing to the

second highest page in terms of mean reward before teleportation.

4.5.1 Sensitivity of optimal solutions to incomplete crawling

Major search engines have spent lots of efforts on crawling the web to discover web pages and the hyperlinks between them. They can thus compute accurately the PageRank. A search engine optimization team may not have such a database available. If one can program a crawler to get a portion of the web graph or download some datasets of reasonable size for free ([web] for instance), these are still incomplete crawlings when compared to the search engine's.

We study here the consequence of considering a portion of the web graph instead of the whole web graph on the strategies computed. In general, one can have no guarantee on the value of the problem because the PageRank could be a lot underestimated if many pages with high PageRank pointing to a given page are not in the truncated crawl. However, as the optimal strategies mainly depend on the mean reward before teleportation v (cf Theorem 4.4), we only need to study the mean reward before teleportation. In this section, we give a bound on the error done with respect to the mean reward before teleportation when considering a truncated crawl instead of Google's crawl. We denote by \tilde{S} , \tilde{P} and \tilde{v} the quantities corresponding to S , P and v in the truncated graph. We define the transition matrix \tilde{S} such that an uncrawled page j is considered to have no link.

For this paragraph, we shall consider the non compensated PageRank. It consists in setting $S^{nc}(j, \cdot) = 0$ for every page j without any link instead of $S(j, \cdot) = z$ and define the non compensated PageRank as $\pi^{nc} = (1 - \alpha)(I - \alpha S^{nc})^{-1}$. Theorem 10 in [Mat04] shows that the non compensated PageRank is proportional to classical PageRank.

Lemma 4.1. *Let v^{nc} be the mean reward before teleportation for the non compensated PageRank, and v the one as in Definition 4.1. Then*

$$\|v - v^{nc}\|_{\infty} \leq \frac{\alpha}{1 - \alpha} zv.$$

Proof. Let $l \in \mathbb{R}$ such that $l_i = 1$ if page i has no outlink and 0 otherwise. Thus, $S^{nc} = S - lz$. We recall that $v = (I_n - \alpha S)^{-1} \bar{r} = \sum_{k \geq 0} \alpha^k S^k \bar{r}$ where $\bar{r}_i = \sum_j P_{i,j} r_{i,j}$. Applying Shermann-Morrisson Formula [Bar51] to $(I - \alpha S + \alpha lz)^{-1}$, we have:

$$v^{nc} = (I - \alpha(S - lz))^{-1} \bar{r} = (I - \alpha S)^{-1} \bar{r} - \alpha (I - \alpha S)^{-1} l \frac{z(I - \alpha S)^{-1} \bar{r}}{1 + \alpha z(I - \alpha S)^{-1} l} = v - (zv)w$$

where $w = \frac{\alpha}{1 + \alpha z(I - \alpha S)^{-1} l} (I - \alpha S)^{-1} l$.

Moreover, as $l \geq 0$ and $\|l\|_{\infty} \leq 1$, $\|w\|_{\infty} \leq \frac{\alpha}{1 - \alpha}$. Finally, $\|v - v^{nc}\|_{\infty} \leq \frac{\alpha}{1 - \alpha} zv$. \square

If only a small portion of the web pages has nonzero instantaneous rewards, then the quantity $(1 - \alpha)zv = \pi \bar{r}$ is the value of the objective and one can have an estimation of its size thanks to the Google toolbar. The mean value of PageRank entries is $\frac{1}{n}$, with n very large, so usually, the absolute value of a given page's PageRank is small compared to unity. In the other hand, the value of the mean reward before teleportation at a controlled page is expected to be of the order of $\|r\|_{\infty}$. Thus the bound given by Lemma 4.1 is rather precise for an "average" website.

We can now prove the following proposition that gives the maximal error done when considering a graph different from Google's when optimizing the PageRank.

Proposition 4.12. *Let I be the set of pages of interest, that is the pages containing or being pointed to by a facultative link. We denote by R the length of a shortest path from a page in I to an uncrawled page. Then, for all i in I ,*

$$|v_i - \tilde{v}_i| \leq 2 \frac{\alpha^{R+1}}{1-\alpha} \|\bar{r}\|_\infty + \frac{\alpha}{1-\alpha} (zv + z\tilde{v})$$

Proof. Using, Lemma 4.1, using S^{nc} and \tilde{S}^{nc} instead of S and \tilde{S} just adds the second summand of the left-hand side.

We have $\forall i \in I, \forall k \leq R, (\tilde{S}^{nc})^k(i, \cdot) = (S^{nc})^k(i, \cdot)$, thus

$$\begin{aligned} |v_i^{nc} - \tilde{v}_i^{nc}| &= \left| \sum_{k \geq 0} \alpha^k e_i((S^{nc})^k - (\tilde{S}^{nc})^k) \bar{r} \right| = \left| \sum_{k \geq R+1} \alpha^k e_i((S^{nc})^k - (\tilde{S}^{nc})^k) \bar{r} \right| \\ &\leq \alpha^{R+1} \|e_i\|_1 \max_k (\|(S^{nc})^k - (\tilde{S}^{nc})^k\|_\infty) \frac{1}{1-\alpha} \|\bar{r}\|_\infty = 2 \frac{\alpha^{R+1}}{1-\alpha} \|\bar{r}\|_\infty . \end{aligned}$$

We have used the property that the infinite norm of a stochastic matrix is 1. \square

This bound is rather large but it is very unlikely to be tight in practice: to reach it, one roughly needs that every page at distance $R+1$ from i is uncrawled and points to page i only. If this situation occurs, this means that there are many pages pointing to our website and that we do not know it. The bound nevertheless shows that the precision of the computation does not depend on the number of pages crawled but on the distance between a controlled page and an uncrawled page.

4.6 PageRank Optimization with coupling constraints

4.6.1 Reduction of the problem with coupling constraints to constrained Markov decision processes

Up to now, we have studied discrete or continuous PageRank Optimization problems but only with local constraints. We consider in this section the following PageRank Optimization problem (4.1) with ergodic (linear in the occupation measure) coupling constraints:

$$\begin{aligned} &\max_{\pi, P} \sum_{i,j} \pi_i P_{i,j} r_{i,j} \text{ st:} \\ &\pi P = \pi, \pi \in \Sigma_n, P_{i,\cdot} \in \mathcal{P}_i, \forall i \in [n] \\ &\sum_{i,j} \pi_i P_{i,j} d_{i,j}^k \leq V^k, \forall k \in K \end{aligned} \tag{4.16}$$

Examples of ergodic coupling constraints are given in Section 4.2.2.

When coupling constraints are present, the previous standard ergodic control model is no longer valid, but we can use instead the theory of *constrained* Markov decision processes. We refer the reader to [Alt99] for more background. In addition to the instantaneous reward r , which is used to define the ergodic functional which is maximized, we now consider a finite family of cost functions $(d^k)_{k \in K}$, together with real constants $(V^k)_{k \in K}$, which will be used to

define the ergodic constraints. The ergodic constrained Markov decision problem consists in finding an admissible control strategy $(\nu_t)_{t \geq 0}$, $\nu_t \in A_{X_t}$, $\forall t \geq 0$, maximizing:

$$\liminf_{T \rightarrow +\infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} r(X_t, \nu_t) \right) \quad (4.17)$$

under the $|K|$ *ergodic constraints*

$$\limsup_{T \rightarrow +\infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} d^k(X_t, \nu_t) \right) \leq V^k, \quad \forall k \in K$$

where the controlled process $(X_t)_{t \geq 0}$ is such that

$$\mathbb{P}(X_{t+1} = j | X_t, \nu_t) = p(j | X_t, \nu_t) .$$

Theorem 4.1 in [Alt99] shows that one can restrict to stationary Markovian strategies and Theorem 4.3 in the same book gives an equivalent formulation of the ergodic constrained Markov decision problem (4.17) as a linear program. When $A_i = \text{extr}(\mathcal{P}_i)$, $r(i, a) = \sum_{j \in [n]} r_{i,j} a_j$, $d^k(i, a) = \sum_{j \in [n]} d_{i,j}^k a_j$ and $p(j | i, a) = a_j$ (see Proposition 4.3), it is easy to see that this linear program is equivalent to:

$$\max_{\rho} \left\{ \sum_{i,j \in [n]} \rho_{i,j} r_{i,j} \text{ st: } \rho \in \mathcal{R} \text{ and } \sum_{i,j \in [n]} \rho_{i,j} d_{i,j}^k \leq V^k, \quad \forall k \in K \right\} \quad (4.18)$$

where \mathcal{R} is the image of $\prod_{i \in [n]} \mathcal{P}_i$ by the correspondence of Proposition 4.1. The set \mathcal{R} is a polyhedron, as soon as every \mathcal{P}_i is a polyhedron (Proposition 4.2).

Following the correspondence discussed in Proposition 4.1, we can see that the linear Problem (4.18) is just the reformulation of Problem (4.16) in terms of occupation measures when we consider total income utility (4.4).

The last result of this section gives a generalization to nonlinear utility functions:

Proposition 4.13. *Assume that the utility function U can be written as $U(P) = W(\rho)$ where W is concave, that the local constraints are convex in P and that the coupling constraints are ergodic. Then, the PageRank Optimization problem (4.16) is equivalent to a concave programming problem in the occupation measure ρ , from which ϵ -solutions can be found in polynomial time.*

Proof. From Proposition 4.2, we know that the set of locally admissible occupation measures is convex. Adding ergodic (linear in the occupation measure) constraints preserves this convexity property. So the whole optimization problem is concave. Finally, Theorem 5.3.1 in [BTN01] states that ϵ -solutions can be found in polynomial time. \square

In particular, the (global) optimality of a given occupation measure can be checked by the first order optimality conditions which are standard in convex analysis.

Remark 4.14. Proposition 4.13 applies in particular if W is a relative entropy utility function, ie. $W(\rho) = - \sum_{i,j \in [n]} \rho_{ij} \log(\rho_{ij} / \mu_{ij})$, where parameters $\mu_{ij} > 0$ (the reference measure) are given.

If we choose to minimize the entropy function on the whole web graph, we recover the TrafficRank algorithm [Tom03]. When we control only some of the hyperlinks whereas the

weights of the others are fixed, the solution of the optimization problem gives the webmaster the weights that she should set to her hyperlinks in order to have an entropic distribution of webservers on her website, interpreted as a fair distribution of webservers.

In the next section, we extend the first order optimality conditions to the formulation in probability transitions, in order to get a characterization of the optimal linking strategies in the constrained PageRank Optimization problem.

4.6.2 Optimality condition

The following shows that the mean reward before teleportation (Definition 4.1) determines the derivative of the utility function. Recall that the *tangent cone* $T_X(x)$ of the set X at point x is the closure of the set of vectors q such that $x + tq \in X$ for t small enough.

Proposition 4.14. *The derivative of the utility function (4.4) is such that for all $Q \in T_{\mathcal{P}}(P)$,*

$$\langle DU(P), Q \rangle = \sum_{i,j} (v_j(P) + r_{i,j}) \pi_i(P) Q_{i,j}$$

where $v(P)$ is the mean reward before teleportation, $\pi(P)$ is the invariant measure of P and $\langle \cdot, \cdot \rangle$ is the standard (Frobenius) scalar product on $n \times n$ matrices.

Proof. We have $U(P) = \sum_{i,j} \pi_i(P) P_{i,j} r_{i,j} = \pi \bar{r}$ and $\pi = \pi P = \pi(\alpha S + (1 - \alpha)ez)$. As $\pi e = 1$, we have an explicit expression for π as function of P : $\pi(P) = (1 - \alpha)z(I_n - P + (1 - \alpha)ez)^{-1}$. The result follows from derivation of $\pi(P)\bar{r}$. We need to derive a product, to derive an inverse ($\langle D(A \mapsto A^{-1}), H \rangle = -A^{-1}HA^{-1}$) and the expression of the mean reward before teleportation $v(P) = (I_n - P + (1 - \alpha)ez)^{-1}\bar{r}$. \square

The next theorem, which involves the mean reward before teleportation, shows that although the continuous constrained PageRank optimization problem is non-convex, the first-order necessary optimality condition is also sufficient.

Theorem 4.5 (Optimality Condition). *Suppose that the sets \mathcal{P}_i defining local constraints are all closed convex sets, that the coupling constraints are given by the ergodic costs functions d^k , $k \in K$ and that the utility function is total income utility. Denote \mathcal{P}^d be the admissible set and $v(P)$ the mean reward before teleportation (Definition 4.1). We introduce the set of saturated constraints $K_{sat} = \{k \in K \mid \sum_{i,j} d_{i,j}^k \pi_i P_{i,j} = V^k\}$ and we denote $D_{i,j}^k = \pi_i d_{i,j}^k + \pi d^k(I - \alpha S)^{-1} e_i P_{i,j}$. Then the tangent cone of \mathcal{P}^d at P is $T_{\mathcal{P}^d}(P) = \left\{ Q \in \prod_{i \in [n]} T_{\mathcal{P}_i}(P_{i,\cdot}) \mid \forall k \in K_{sat}, \langle D^k, Q \rangle \leq 0 \right\}$ and $P^* \in \mathcal{P}^d$ is the optimum of the continuous PageRank Optimization problem (4.1) with ergodic coupling constraints if and only if:*

$$\forall Q \in T_{\mathcal{P}^d}(P^*), \quad \sum_{i,j \in [n]} \pi_i (v_j(P^*) + r_{i,j}) Q_{i,j} \leq 0$$

Proof. Let us consider the birational change of variables of Proposition 4.1. As all the occupation measures considered are irreducible, its Jacobian is invertible at any admissible point. Thus, we can use the results of Section 6.C in [RW97]. Denote $\mathcal{P} = \prod_{i \in [n]} \mathcal{P}_i$, with tangent cone $T_{\mathcal{P}}(P) = \prod_{i \in [n]} T_{\mathcal{P}_i}(P_{i,\cdot})$, and $\mathcal{R} = f^{-1}(\mathcal{P})$. We have $T_{\mathcal{R}^d}(\rho) = \left\{ \sigma \in T_{\mathcal{R}}(\rho) \mid \forall k \in K_{sat}, \langle d^k, \sigma \rangle \leq 0 \right\}$ and $T_{\mathcal{P}^d}(P) = \left\{ Q \in \mathbb{R}^{n \times n} \mid \nabla f^{-1} Q \in T_{\mathcal{R}^d}(f^{-1}(P)) \right\}$.

$\nabla f^{-1}Q \in \mathbb{T}_{\mathcal{R}^d}(f^{-1}(P))$ first means that $\nabla f^{-1}Q \in \mathbb{T}_{\mathcal{R}}(f^{-1}(P))$ which can also be written as $\nabla f \nabla f^{-1}Q = Q \in \mathbb{T}_{\mathcal{P}}(P)$. The second condition is $\forall k \in K_{sat}, \langle d^k, \nabla f^{-1}Q \rangle \leq 0$. As $(f^{-1}(P))_{i,j} = \rho_{i,j} = \pi_i P_{i,j}$, we have $(\nabla f^{-1}Q)_{i,j} = \sum_{k,l} Q_{k,l} (P_{k,l} \frac{\partial \pi_k}{\partial P_{i,j}} + \pi_k \delta_{ik} \delta_{jl})$. Thanks to the expression the derivative of the utility function and of $\frac{\partial \pi_k}{\partial P_{i,j}} = \pi_i e_j (I - \alpha S)^{-1} e_k$ both given in Proposition 4.14, we get the expression stated in the theorem.

By Proposition 4.13, the PageRank optimization problem is a concave programming problem in ρ and so, the first order (Euler) optimality condition guarantees the global optimality of a given measure. Thus, every stationary point for the continuous PageRank Optimization problem is a global maximum when written in transition probabilities also. \square

4.6.3 A Lagrangian relaxation scheme to handle coupling constraints between pages

The PageRank Optimization Problem with "ergodic" coupling constraints (4.16) may be solved by off the shelf simplex or interior points solvers. However, such general purpose solvers may be too slow, or too memory consuming, to solve the largest web instances.

The following proposition yields an algorithm that decouples the computation effort due to complexity of the graph and due to coupling constraints.

Proposition 4.15. *The PageRank Optimization problem with K "ergodic" coupling constraints (4.16) can be solved by a Lagrangian relaxation scheme, in which the dual function and one of its subgradient*

$$\theta(\lambda) = \max_{\rho \in \mathcal{R}} \langle r, \rho \rangle - \sum_{k \in K} \lambda_k (\langle d^k, \rho \rangle - V^k)$$

$$\frac{\partial \theta}{\partial \lambda_k}(\lambda) = \langle d^k, \rho^*(\lambda) \rangle - V^k$$

are evaluated by dynamic programming and $\rho^*(\lambda)$ is a maximizer of the expression defining $\theta(\lambda)$.

Proof. This is a simple application of Lagrange multipliers theory, see [Lem01] Theorem 21 and Remark 33 for instance. Here we relax the coupling constraints in the problem written with occupation measures (4.18). We solve the dual problem, namely we minimize the dual function θ on \mathbb{R}_+^K . The value of this dual problem is the same as the value of the constrained primal problem and we can get a solution of the primal problem since there is no duality gap. \square

We have implemented a bundle high level algorithm, in which the dual function is evaluated at each step by running a value iteration algorithm, for a problem with modified reward. By comparison with the unconstrained case, the execution time is essentially multiplied by the number of iterations of the bundle algorithm.

4.6.4 A heuristic to solve the discrete problem with coupling constraints

In this section, we give an upper bound for the discrete PageRank optimization problem with coupling constraints and a heuristic for the PageRank optimization problem with mutual exclusion constraints.

Instead of solving these difficult problems, we consider the continuous problem with the polytopes of uniform transition measures as local admissible sets, i.e. we relax the discrete pattern. Thus by the Lagrangian scheme of Proposition 4.15, we get an upper bound on the optimal objective and we have a lower bound for any admissible discrete transition matrix (see Section 4.7.3).

When it comes to the PageRank optimization problem with mutual exclusion constraints, we have to deal with combinatorial constraints. However, the problem can be written as a mixed linear program:

$$\begin{aligned} & \max_{\rho, x} \sum_{i,j} r_{i,j} \rho_{i,j} \\ & \rho_{i,j} - \underline{\rho}_{i,j} \leq M_{i,j} x_{i,j}, \text{ for all constrained link } (i, j) \\ & x_{i_1, j_1} + x_{i_2, j_2} \leq 1, \text{ for all exclusion constraint between } (i_1, j_1) \text{ and } (i_2, j_2) \\ & \rho \in \mathcal{R} \quad x \in \{0, 1\}^K \end{aligned} \tag{4.19}$$

where \mathcal{R} is the set of occupation measures satisfying the local constraints, $\underline{\rho}_{i,j} = (1 - \alpha) z_j \sum_l \rho_{i,l}$ is the value of the occupation measure where there is no link, $M_{i,j}$ is an upper bound on $\rho_{i,j} - \underline{\rho}_{i,j}$ given for instance by a PageRank optimization problem with local constraints and $x_{i,j}$ is a binary variable that forces the absence of the link from page i to page j when set to 1.

We consider the continuous problem with the polytopes of uniform transition measures as local admissible sets and $x \in [0, 1]^K$, i.e. we relax the discrete pattern. Thus we get an upper bound on the optimal objective. To compute it, we remark that at fixed x , the problem is a PageRank optimization problem with linear coupling constraints and thus is tractable by Proposition 4.15. The remaining function of x is a nondifferentiable concave function with as many variables as pages with a mutual exclusion constraint. Hence it is tractable by convex programming techniques.

The heuristic to get feasible points is based on the following observation: when x is binary, then the coupling constraints in (4.19) either reduce to forbidden links constraints ($x_{i,j} = 0$) or are useless ($x_{i,j} = 1$). This means that there are no more coupling constraints and finding a discrete solution is easy. In order to get a discrete point from the fractional solution given by the convex program, we select the biggest coordinate of x , we set it to 1 and we set the coordinates with which it has an exclusion constraint to 0. We then start back the process with the next fractional coordinate until we have a binary vector. The result is not optimal but provides a lower bound that can be compared with the upper bound.

It may also be possible to design a branch and bound algorithm to solve the problem exactly thanks to the bounds found.

Remark 4.15. If there are local exclusion constraints, an alternative formulation is to set directly the constraint $S_{i,j} + S_{i,k} \leq S_{i,j_0}$ where j_0 is an obligatory link. Such constraints lead to a polytope where all extreme points are uniform transition probabilities as soon as the graph defined such that there is an arc between the hyperlinks in competition is bipartite.

4.7 Experimental results

We have tried our algorithms on a 2006 crawl on eight New Zealand Universities available at [Pro06]. There are 413,639 nodes and 2,668,244 links in the graph. The controlled set we

have chosen is the set of pages containing "maori" in their url. There are 1292 of them. We launched the experiments in a sequential manner on a personal computer with Intel Xeon CPU at 2.98 Ghz and wrote the code in Scilab language.

4.7.1 Continuous problem with local constraints only

Assume that the webmasters controlling the pages containing "maori" in their url cooperate and agree to change at most 20% of the links' weight to improve the PageRank, being understood that self-links are forbidden (skeleton constraint, see Section 4.2.2). The algorithm launched on the optimization of the sum of the PageRanks of the controlled pages (calculated with respect to the crawled graph only, not by the world wide graph considered by Google) ran 27 seconds.

The optimal strategy returned is that every controlled page except itself should link with 20% weight to the page `maori-oteha.massey.ac.nz/te_waka.htm`. That page should link to the page `maori-oteha.massey.ac.nz/tewaka/about.htm`. The sum of PageRank values goes from 0.0057 to 0.0085.

Hence, by uniting, this team of webmasters would improve the sum of their PageRank scores by 49%. Remark that all the pages point to the same page (except itself because self-links are forbidden). The two best pages to point to are in fact part of a "dead end" of the web graph containing only pages with maximal reward. A random surfer can only escape from this area of the graph by teleporting, which makes the mean reward before teleportation maximal.

4.7.2 Discrete problem

On the same data set, we have considered the discrete optimization problem. The set of obligatory links is the initial set of links. We have then selected 2,319,174 facultative links on the set of controlled pages of preceding section.

Execution time took 81 seconds with the polyhedral approach of Section 4.4.2 (60 iterations). We compared our algorithm with the adaptation of the graph augmentation approach of [CJB10] to total utility presented in Section 4.4.3: this algorithm took 460 seconds (350 iterations) for the same precision. The optimal strategy is to add no link that goes out of the website but get the internal link structure a lot denser. From 12,288 internal links, the optimal strategy is to add 962,873 internal links. Finally, 98.2% of the links are internal links and there is a mean number of links per page of 770. The sum of PageRank values jumps from 0.0057 to 0.0148.

Here, as the weights of the links cannot be changed, the webmaster can hardly force websurfers to go to dead ends. But she can add so many links that websurfers get lost in the labyrinth of her site and do not find the outlinks, even if they were obligatory.

4.7.3 Coupling linear constraints

As we have seen in the preceding experiments, optimizing the PageRank with too much freedom may lead to an awkward and unpleasant website. So we would like to solve the discrete optimization problem of the preceding section with additional design constraints. We require that each visitor coming on one of the pages of the team has a probability to leave the set of pages of the team on next step of 40% (coupling conditional probability constraint, see Section 4.2.2). This guarantees that websurfers will not be led to dead ends. We also require

that the sum of PageRank values of the home pages of the 10 universities considered remains at least equal to their initial value after the optimization (effective frequency constraint). Finally, we limit the number of links added by page to 20 (local constraint discussed in Remark 4.9).

In the case of constrained Markov decision processes, optimal strategies are usually randomized strategies. This means that the theory cannot directly deal with discrete action sets. Instead, we consider the continuous problem with the polytopes of uniform transition measures as local admissible sets, i.e. we relax the discrete pattern. Thus by the Lagrangian scheme of Proposition 4.15, we get an upper bound on the optimal objective and we have a lower bound for any admissible discrete transition matrix.

The initial value is 0.0057 and the Lagrangian relaxation scheme gives an upper bound of 0.00742. Computation took 500 s (8 high level iterations). During the course of the Lagrangian relaxation scheme, all intermediate solutions are discrete and three of them satisfied the coupling constraints. The best of them corresponds to a sum of PageRanks of 0.00739 (improvement: 30%). Thus the duality gap is at most 0.4%. In general, the intermediate discrete solutions need not satisfy the coupling constraints and getting an admissible discrete solution may be difficult.

The discrete transition matrix found suggests to add 15,324 internal links but also 2,851 external links. Thanks to the limit on the number of links added, fewer hyperlinks are added than in Section 4.7.2. Moreover, thanks to the coupling constraints, external links are added too, which improves the quality of the website.

Conclusion

We have presented in this chapter a general framework to study the optimization of PageRank. Our results apply to a continuous problem where the webmaster can choose the weights of the hyperlinks on her pages and to the discrete problem in which a binary decision must be taken to decide whether a link is present. We have shown that the Discrete PageRank Optimization problem without coupling constraints can be solved by reduction to a concisely described relaxed continuous problem. We also showed that the continuous PageRank optimization problem is polynomial time solvable, even with coupling constraints.

We gave scalable algorithms which rely on an ergodic control model and on dynamic programming techniques. The first one, which applies to problems with local design constraints, is a fixed point scheme whose convergence rate shows that optimizing PageRank is not much more complicated than computing it. The second algorithm, which handles coupling constraints, is still efficient when the number of coupling constraints remains small.

We have seen that the mean reward before teleportation gives a total order of preference in pointing to a page or another. This implies that pages high in this order concentrate many inlinks from controlled pages. This is a rather degenerate strategy when we keep in mind that a web site should convey information. Nevertheless, the model allows one to address more complex problems, for instance with coupling constraints, in order to get less trivial optimal linking strategies.

This work may be useful to understand link spamming, to price Internet advertisements or, by changing the objective function, to design web sites with other goals like fairness or usefulness.

CHAPTER 5

PageRank optimization applied to spam detection

5.1 Introduction

From the early days of search engines, some webmasters have tried to get their web pages overranked thanks to malicious manipulations. For instance, adding many keywords on a page is a classical way to make search engines consider a page relevant to many queries. With the advent of link-based algorithms, spammers have developed new strategies, called link-spamming [GGM05b], that intend to give some target page a high score. For instance, Gyöngyi and Garcia-Molina [GGM05a] showed various linking strategies that improve the PageRank score of a page. They justified the presence of link farms with patterns with every page linking to one single page. Baeza-Yates, Castillo and López [BYCL05] also showed that making collusions is a good way to improve PageRank.

In order to fight such malicious manipulations that deteriorate search engines' results and deceive web surfers, various techniques have been developed. We refer to [CD10] for a detailed survey of this subject. Each spam detection algorithm is focused on a particular aspect of spam pages. Content analysis (see [NNMF06] for instance) is the main tool to detect deceiving keywords. Some simple heuristics [AAD08] may be enough to detect the most coarse link-spam techniques, but more evolved graph algorithms like clique detection [STKA07], SpamRank [BCSU05] or Truncated PageRank [BL06] have also been developed to fight link-

spamming. As web spammers adapt themselves to detection algorithms, machine learning techniques [GS07] try to discover actual predominant spam strategies and to adapt to its evolutions. Another direction of research concerns the propagation of trust through the web graph with the TrustRank algorithm [GGMP04] and its variants or the propagation of distrust through a reversed web graph with the AntiTrustRank algorithm [KR06].

In this chapter, we develop a new link spam detection and PageRank demotion algorithm called MaxRank. Like in [GGMP04, KR06], we start with a seed of hand-picked trusted and spam pages. We define the MaxRank of a page as the frequency of visit of this page by a random surfer minimizing an average cost per time unit. On a given page, the random surfer selects a set of hyperlinks and clicks with uniform probability on any of these hyperlinks. The cost function penalizes spam pages and hyperlink removals. The goal is to determine an optimal hyperlink deletion policy. The features of MaxRank are based on PageRank optimization [AL06, MV06, dKNvD08, IT09, CJB10] and more particularly on the results of Chapter 4. Those works have shown that the problem of optimizing the PageRank of a set of pages by controlling some hyperlinks can be solved by Markov Decision process algorithms. There, the optimization of PageRank was thought from a webmaster’s point of view whereas here, we take the search engine’s point of view. We show that the Markov Decision Process defining MaxRank is solvable in polynomial time (Proposition 5.2), because the polytopes of occupation measures admit efficient (polynomial time) separation oracles. The invariant measure of the Markov Decision Process, the MaxRank vector, is interpreted as a modified PageRank vector, used to sort web pages instead of the usual PageRank vector. The solution of the ergodic dynamic programming equation, called the bias vector, is unique up to an additive constant. We show that it can be interpreted as a measure of the “spamicity” of each page, used to detect spam pages.

We give a scalable algorithm for MaxRank computation that allowed us to perform numerical experiments on the WEBSpAM-UK2007 dataset [web07]. We show that our algorithm outperforms both TrustRank and AntiTrustRank for spam and nonspam page detection. As an example, on the WEBSpAM-UK2007 dataset, for a recall of 0.8 in the spam detection problem, MaxRank has a precision of 0.87 while TrustRank has a precision of 0.30 and AntiTrustRank a precision of 0.13.

5.2 The MaxRank algorithm

In this section, we define the MaxRank algorithm. It is based on our earlier works on PageRank optimization. In Chapter 4, we considered the problem of optimizing the PageRank of a given website from a webmaster’s point of view, that is with some controlled hyperlinks and design constraints. Here, we take the search engine’s point of view. Hence, for every hyperlink of the web, we can choose to take it into account or not: our goal is to forget spam links while letting trusted links active in the determination of the ranking.

As in TrustRank [GGMP04] and AntiTrustRank [KR06], we start with a seed of trusted pages and known spam pages. The basic idea is to minimize the sum of PageRank scores of spam pages and maximize the sum of PageRank scores of nonspam pages, by allowing us to remove some hyperlinks when computing the PageRank. However, if we do so, the optimal strategy simply consists in isolating already known spam pages from trusted pages: there is then no hope to detect other spam and nonspam pages. Thus, we add a penalty when a hyperlink is removed, so that “spamicity” can still propagate through (removed or

not removed) hyperlinks. Finally, we also control the teleportation vector in order to penalize further pages that we suspect to be spam pages.

We model this by a controlled random walk on the web graph, in which the hyperlinks can be removed. Each time the random surfer goes to a spam page, he gets a positive cost, each time he goes to a trusted page, he gets a negative cost. When the status of the page is unknown, no cost incurs. In addition to this a priori cost, he gets a penalty for each hyperlink removed. Like for PageRank, the random surfer teleports with probability α at every time step; however, in this framework, he chooses the set of pages to which he wants to teleport.

Let \mathcal{F}_x be the set of pages pointed by x in the original graph and D_x be the degree of x . An action consists in determining $J \subseteq \mathcal{F}_x$, the set of hyperlinks that remain, and $I \subseteq [n]$, the set of pages to which the surfer may teleport. We shall restrict I to have a cardinality equal to $N \leq n$. Then, following (2.2), the probability of transition from page x to page y is

$$p(y|x, I, J) = \alpha \nu_y(I, J) + (1 - \alpha) z_y(I)$$

where the teleportation vector and the hyperlink click probability distribution are given by

$$z_y(I) = \begin{cases} |I|^{-1} & \text{if } y \in I \\ 0 & \text{otherwise} \end{cases}$$

$$\nu_y(I, J) = \begin{cases} z_y(I) & \text{if } |J| = \emptyset \\ |J|^{-1} & \text{if } y \in J \\ 0 & \text{otherwise} \end{cases}$$

The cost at page x is given by

$$c(x, I, J) = c'_x + \gamma \frac{D_x - |J|}{D_x} .$$

c'_x is the a priori cost of page x . This a priori cost should be positive for known spam pages and negative for trusted pages. D_x is the degree of x in the web graph and $\gamma > 0$ is a penalty factor. The penalty $\gamma \frac{D_x - |J|}{D_x}$ is proportional to the number of pages removed.

We study the following ergodic control problem:

$$\inf_{(I_t)_{t \geq 0}, (J_t)_{t \geq 0}} \limsup_{T \rightarrow +\infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} c(X_t, I_t, J_t) \right) , \quad (5.1)$$

where an admissible control consists in selecting, at each time step t , a subset of pages $I_t \subseteq [n]$ with $|I_t| = N$ to which teleportation is permitted, and a subset $J_t \subseteq \mathcal{F}_{X_t}$ of the set of hyperlinks in the currently visited page X_t .

The following proposition gives an alternative formulation of Problem (5.1) that we will then show to be well-described.

Proposition 5.1. *Fix $N \in \mathbb{N}$ and let*

$$Z = \left\{ z \in \mathbb{R}^n \mid \sum_{i \in [n]} z_i = 1, 0 \leq z_i \leq \frac{1}{N} \right\} .$$

Let \mathcal{F}_x be the set of pages pointed by x in the original graph and D_x be the degree of x . Let \mathcal{P}_x be the polyhedron defined as the set of vectors $(\sigma, \nu) \in \mathbb{R}^{D_x+1} \times \mathbb{R}^n$ such that there exists $w \in \mathbb{R}^{(D_x+1) \times n}$ verifying

$$\sum_{d=0}^{D_x} \sigma^d = 1 \quad (5.2a)$$

$$\sigma^d \geq 0, \quad \forall d \in \{0, \dots, D_x\} \quad (5.2b)$$

$$\nu_j = \sum_{d=0}^{D_x} w_j^d, \quad \forall j \in [n] \quad (5.2c)$$

$$\sum_{j \in [n]} w_j^d = \sigma^d, \quad \forall d \in \{0, \dots, D_x\} \quad (5.2d)$$

$$0 \leq w_j^0 \leq \frac{\sigma^0}{N_z}, \quad \forall j \in [n] \quad (5.2e)$$

$$w_j^d = 0, \quad \forall j \notin \mathcal{F}_x, \forall d \in \{1, \dots, D_x\} \quad (5.2f)$$

$$0 \leq w_j^d \leq \frac{\sigma^d}{d}, \quad \forall j \in \mathcal{F}_x, \forall d \in \{1, \dots, D_x\} \quad (5.2g)$$

Then Problem 5.1 is equivalent to the following ergodic control problem

$$\inf_{\sigma, \nu, z} \limsup_{T \rightarrow +\infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} \tilde{c}(X_t, \sigma_t, \nu_t, z_t) \right), \quad (5.3)$$

where the cost is defined as

$$\tilde{c}(x, \sigma, \nu, z) = c'_x + \gamma \frac{D_x - \sum_{d=0}^{D_x} d\sigma^d}{D_x}$$

and the transitions are

$$\tilde{p}(y|x, \sigma, \nu, z) = \alpha \nu_y + (1 - \alpha) z_y.$$

The admissible controls verify for all t , $(\sigma_t, \nu_t) \in \text{extr}(\mathcal{P}_{X_t})$ (the set of extreme point of the polytope) and $z_t \in \text{extr}(Z)$.

Indeed, to each action (σ, ν, z) of Problem (5.3) corresponds a unique action I, J of Problem (5.1) and vice versa. Moreover, the respective transitions and costs are equal.

Proof. Fix a page x in $[n]$. The extreme points of Z are the vectors of \mathbb{R}^n with N coordinates equal to $\frac{1}{N}$ and the other ones equal to 0. Hence $z(I) \in \text{extr}(Z)$ and for each extreme point z' of Z there exists $I \in [n]$ such that $|I| = N$ and $z' = z(I)$. We shall also describe the set of extreme points of \mathcal{P}_x .

From the theory of disjunctive linear programming [Bal98], we can see that the polytope $\mathcal{K} = \{\nu \mid (\sigma, \nu) \in \mathcal{P}_x\}$ is the convex hull of the union of $D_x + 1$ polytopes that we will denote K_d , $d \in \{0, 1 \dots D_x\}$. If $d = 0$, then $K_0 = Z$. If $d > 0$, $K_d = \{\nu \in \mathbb{R}^n \mid \sum_{j \in [n]} \nu_j = 1, 0 \leq \nu_j \leq \frac{1}{d}, \forall j \in \mathcal{F}_x, \nu_j = 0, \forall j \notin \mathcal{F}_x\}$.

Let (σ, ν) be an extreme point of \mathcal{P}_x . By Corollary 2.1.2-ii) in [Bal98], there exists d^* such that $\sigma^{d^*} = 1$ and ν is an extreme point of \mathcal{K} . As $\sigma^{d^*} = 1$ and $\sigma_d = 0$ for $d \neq d^*$, we conclude that ν is also an extreme point of $K_{d^*} = \mathcal{P}_x \cap \{\sigma \mid \sigma^{d^*} = 1\}$. If $d^* = 0$, $\nu \in \text{extr}(Z)$ and if

$d^* > 0$, the extreme points of K_{d^*} correspond exactly to the vectors with d^* coordinates in \mathcal{F}_x equal to $\frac{1}{d^*}$ and the other ones equal to 0. Hence there exists $I \subseteq [n]$ and $J \subseteq \mathcal{F}_x$ such that $|I| = N$, $|J| = d^*$ and $\nu = \nu(I, J)$.

Conversely, fix I and J and let $\sigma(J)$ be such that $\sigma^d(J) = 1$ if and only if $d = |J|$. Then $(\sigma(J), \nu(I, J))$ is an extreme point of \mathcal{P}_x by Corollary 2.1.2-i) in [Bal98].

Finally the costs are the same since $|J| = \sum_{d=0}^{D_x} d\sigma^d$. \square

Proposition 5.2. *If α , γ and c'_i , $i \in [n]$ are rational numbers, then the ergodic control problem (5.3) is the average cost infinite horizon problem for a well described Markov decision process and it is polynomial time solvable.*

Proof. Clearly, the process described is a Markov decision process. As the polytopes \mathcal{P}_i , $i \in [n]$ and Z are described by a polynomial number of inequalities with at most $n + 1$ terms in each, they are well described. Indeed, the separation oracle consisting simply in testing each inequality terminates in polynomial time. The cost and transitions are linear functions on those polytopes with rational coefficients since α and c'_i , $i \in [n]$ are rational numbers. Thus the Markov decision process is well described. By Theorem 3 in [FABG13], Problem (5.3) is thus solvable in polynomial time. \square

Proposition 5.3. *The dynamic programming equation*

$$v_i + \lambda = \min_{(\sigma, \nu) \in \mathcal{P}_i, z \in Z} c'_i + \gamma \frac{D_i - \sum_{d=0}^{D_i} d\sigma^d}{D_i} + \sum_{j \in [n]} (\alpha \nu_j + (1 - \alpha) z_j) v_j, \quad \forall i \in [n] \quad (5.4)$$

has a solution $v \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$. The constant λ is unique and is the value of problem (5.1). An optimal strategy is obtained by selecting for each state i , $(\sigma, \nu) \in \mathcal{P}_i$ and $z \in Z$ maximizing Equation (5.4). The function v is called the bias.

Proof. Theorem 8.4.3 in [Put94] applied to the unichain ergodic control problem (5.3) implies the result of the proposition but with \mathcal{P}_i replaced by $\text{extr}(\mathcal{P}_i)$. But as the expression which is maximized is affine, using \mathcal{P}_i or $\text{extr}(\mathcal{P}_i)$ yields the same solution. Proposition 5.1 gives the equivalence between (5.1) and (5.3) \square

Proposition 5.4. *Let T be the dynamic programming operator $\mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by*

$$T_i(v) = \min_{(\sigma, \nu) \in \mathcal{P}_i} c'_i + \gamma \frac{D_i - \sum_{d=0}^{D_i} d\sigma^d}{D_i} + \alpha \sum_{j \in [n]} \nu_j v_j, \quad \forall i \in [n].$$

The map T is α -contracting in the sup norm and its fixed point v , which is unique, is such that $(v, (1 - \alpha) \min_{z \in Z} z \cdot v)$ is solution of the ergodic dynamic programming equation (5.4).

Proof. The set $\{\nu \text{ st: } (\sigma, \nu) \in \mathcal{P}_i\}$ is a set of probability measures so $\lambda \in \mathbb{R} \Rightarrow T(v + \lambda) = T(v) + \alpha \lambda$ and $v \geq w \Rightarrow T(v) \geq T(w)$. This implies that T is α -contracting. Let v be its fixed point. For all $i \in [n]$,

$$v_i + (1 - \alpha) \min_{z \in Z} z \cdot v = \min_{(\sigma, \nu) \in \mathcal{P}_i, z \in Z} c'_i + \gamma \frac{D_i - \sum_{d=0}^{D_i} d\sigma^d}{D_i} + \sum_{j \in [n]} \alpha \nu_j v_j + (1 - \alpha) z_j v_j$$

We get equation (5.4) with constant $(1 - \alpha) \min_{z \in Z} z \cdot v$. \square

Remark 5.1. The dynamic programming operator T is a monotone and additively sub-homogeneous map. Hence, v can be seen as its eigenvector corresponding to the eigenvalue 0. The invariant measure is here a kind of “left eigenvector” for this nonlinear map. Hence we see that the MaxRank algorithm also belongs to the class of Perron rankings (see Section 2.3).

We can then solve the dynamic programming equation (5.4) and so the ergodic control problem (5.1) by value iteration (outer loop of Algorithm 5.1).

The algorithm starts with an initial potential function v , scans repeatedly the pages and updates v_i when i is the current page according to $v_i \leftarrow T_i(v)$ until convergence is reached. Then $(v, (1 - \alpha) \min_{z \in Z} zv)$ is solution of the ergodic dynamic programming equation (5.4) and an optimal linkage strategy is recovered by selecting the maximizing (σ, ν) at each page. An optimal teleportation vector is recovered by selecting a maximizing z in $\min_{z \in Z} zv$.

Thanks to the damping factor α , the iteration can be seen to be α -contracting if the pages are scanned in a cyclic order. Thus the algorithm converges in a number of steps independent of the dimension of the web graph.

For the evaluation of the dynamic programming operator T at a page i (inner for-loop of Algorithm 5.1), we remark that for a fixed $0 - 1$ valued σ , that is for a fixed number of removed hyperlinks, the operator T_i maximizes a linear function on a hypercube, which reduces essentially to a sort. Any extreme point (σ, ν) of \mathcal{P}_i necessarily verifies that σ is $0 - 1$ valued. Thus we just need to choose the best value of σ among the $D_i + 1$ possibilities.

Algorithm 5.1 MaxRank algorithm

- 1: Initialization: $v \in \mathbb{R}^n$
 - 2: **while** $\|v - T(v)\|_\infty \geq \epsilon$ **do**
 - 3: Sort $(v_l)_{l \in [n]}$ in increasing order and let $\phi : [n] \rightarrow [n]$ be the sort function so that $v_{\phi(1)} \leq \dots \leq v_{\phi(n)}$.
 - 4: $\lambda \leftarrow \frac{1-\alpha}{N} \sum_{j=1}^N v_{\phi(j)}$
 - 5: **for** i from 1 to n **do**
 - 6: $w_i^0 \leftarrow c'_i + \gamma + \frac{\alpha}{1-\alpha} \lambda$
 - 7: Sort $(v_j)_{j \in \mathcal{F}_i}$ in increasing order and let $\psi : \mathcal{F}_i \rightarrow \{1, \dots, |\mathcal{F}_i|\}$ be the function such that $v_{\psi(1)} \leq \dots \leq v_{\psi(|\mathcal{F}_i|)}$.
 - 8: **for** d from 1 to D_i **do**
 - 9: $w_i^d \leftarrow c'_i + \gamma \frac{D_i - d}{D_i} + \frac{\alpha}{d} \sum_{j=1}^d v_{\psi(j)}$
 - 10: **end for**
 - 11: $T_i(v) = \min_{d \in \{0, 1, \dots, D_i\}} w_i^d$
 - 12: **end for**
 - 13: $v \leftarrow T(v)$
 - 14: **end while**
-

This very efficient algorithm is highly scalable: we used it for our experimental results on a large size dataset (Section 5.3).

The following proposition shows that if γ is too big, then the optimal link removal strategy is trivial. It also gives an interpretation of the bias vector in terms of number of visits of spam pages.

Proposition 5.5. *If $\gamma > \frac{2\alpha}{1-\alpha} \|c\|_\infty$, then no link should be removed. Moreover if in addition, $c'_i = 1$ when i is a spam page and 0 otherwise, then the i th coordinate of the fixed point of*

operator T in Proposition 5.4 (the bias vector) is equal to the expected mean number of spam pages visited before teleportation when starting a walk from page i .

Proof. Proposition 5.4 gives a normalization of the bias vector such that it cannot take any value greater than $\frac{\|c\|_\infty}{1-\alpha}$, as it is the case for PageRank optimization [dKNvD08, FABG13].

Now fix $i \in [n]$. Let $\nu^0 = \nu(\mathcal{F}_i)$ and $\nu = \nu(I, J)$ for $J \subset \mathcal{F}_i$. If $J = \emptyset$, then

$$\alpha|\nu v - \nu^0 v| \leq \alpha|\nu v| + \alpha|\nu^0 v| \leq \frac{2\alpha}{1-\alpha} \|c\|_\infty \frac{D_i - |\emptyset|}{D_i}.$$

If $J \neq \emptyset$, then

$$\begin{aligned} \alpha|\nu v - \nu^0 v| &\leq \alpha \left(\frac{1}{|J|} - \frac{1}{D_i} \right) \sum_{j \in J} v_j + \alpha \frac{1}{D_i} \sum_{j \in \mathcal{F}_i} v_j \\ &\leq \frac{D_i - |J|}{|J| D_i} \frac{\alpha \|c\|_\infty}{1-\alpha} + \frac{D_i - |J|}{D_i} \frac{\alpha \|c\|_\infty}{1-\alpha} \leq \gamma \frac{D_i - |J|}{D_i}. \end{aligned}$$

This proves that choosing $J = \mathcal{F}_i$ is always the best strategy when $\gamma > \frac{2\alpha}{1-\alpha} \|c\|_\infty$.

When no link is removed and c' is defined as in the proposition, we are in the framework of [dKNvD08], where it is shown that the i th coordinate of the fixed points of the operator T is equal to the expected mean number of visits before teleportation when starting a walk from page i . \square

5.3 Spam detection and PageRank demotion

We performed numerical experiments on the WEBSpAM-UK2007 dataset [web07]. This is a crawl of the .uk websites with $n = 105,896,555$ pages performed in 2007, associated with lists of hosts classified as spam, nonspam and borderline. There is a training dataset for the setting of the algorithm and a test dataset to test the performance of the algorithm.

We took $\gamma = 4$, $\alpha = 0.85$, $N = 0.89n$, $c'_i = 1$ if i is a spam page of the training dataset, $c'_i = -0.2$ if i is a nonspam page of the training dataset and $c'_i = 0$ otherwise. Then we obtained the MaxRank score and the associated bias vector. We also computed TrustRank and AntiTrustRank with the training dataset as the seed sets. We used the Webgraph framework [BV04], so that we could manage the computation on a personal computer with four Intel Xeon CPUs at 2.98 GHz and 8 GB RAM. We coded the algorithm in a parallel fashion thanks to the OpenMP library. Computation took around 6 minutes for each evaluation of the dynamic programming operator of Proposition 5.4 and 6 hours for 60 such iterations (precision on the objective $\alpha^{60} \leq 6.10^{-5}$). By comparison, PageRank computation with the same precision required 1.3 hour on the same computer, which is of the same order of magnitude.

Figure 5.1 gives the values taken by the bias vector. Figure 5.2 compares the precision and recall of PageRank, TrustRank, AntiTrustRank and MaxRank bias for spam or non spam detection. Precision and recall are the usual measures of the quality of an information retrieval algorithm [BYRN99]. Precision is the probability that a randomly selected retrieved document is relevant. Recall is the probability that a randomly selected relevant document is retrieved in a search. These values were obtained using the training and the test sets. Figure 5.3 compares TrustRank and MaxRank scores for PageRank demotion. It shows that MaxRank is a good candidate for a spam-resistant version of PageRank.




Figure 5.1: Recognizing spam with the bias vector. Top: the values of the bias for all pages. Bottom: zoom on bias values near 0. Pages are sorted by growing bias value. Spam pages of the training set have a large positive bias value, non spam pages of the training set have a negative bias value. Pages in between describe a “continuum” of values and the separation between pages considered spam or not is arbitrarily set.




Figure 5.2: Top: Precision as a function of recall for spam detection. Bottom: Precision as a function of recall for detection of non spam pages. We present the result for four algorithms: PageRank (dotted line below), TrustRank (dashed line), AntiTrustRank (dash-dotted line) and MaxRank bias (solid line). Given a vector of scores, the various pairs of precision-recall values are obtained for various thresholds defining the discrimination between pages considered spam and non spam. TrustRank and AntiTrustRank have a precision of 1 on their seed set, which represent around 70% of the total test set. But out of their training set, their precision decreases quickly. MaxRank, on the other hand, remains precise even out of its training set.




Figure 5.3: Top: PageRank demotion by TrustRank. We show the ratio of TrustRank over PageRank. Pages are sorted by growing promotion. The demotion is very coarse, since there are many pages with a very small TrustRank score, when compared to their PageRank score. Bottom: PageRank demotion by MaxRank. The demotion is finer, most of the pages have a nearly unchanged score, up to a multiplicative constant.

CHAPTER 6

Convergence of Tomlin's HOTS algorithm

6.1 Introduction

In this chapter, we focus on an algorithm proposed by Tomlin in [Tom03] for the ranking of web pages, called HOTS. It may also be used for other purposes like the ranking of sport teams [Gov08]. Like PageRank [BP98], HITS [Kle99] and SALSA [LM00], HOTS uses the hyperlink structure of the web (see also [LM05b, LM06] for surveys on link-based ranking algorithms). This structure is summarized in the web graph, which is the digraph with a node for each web page and an arc between pages i and j if there is a hyperlink from page i to page j .

The HOTS vector, used to rank web pages, is the vector of the exponentials of the dual variables of an optimal flow problem. The flow represents an optimal distribution of web surfers on the web graph in the sense of entropy maximization. The dual variable, one by page, is interpreted as the “temperature” of the page, the hotter a page the better. In the case of the PageRank, the flow of websurfers is determined by the uniform transition probability of following one hyperlink in the current page. This transition rule is in fact arbitrary. The HOTS model assumes that the web surfers choose the hyperlink to follow by maximizing the entropy of the flow. Hence, the HOTS problem is a continuous PageRank optimization problem where the weights of all the hyperlinks of the web are controlled and the objective function is the entropy of the occupation measure (The occupation measure as defined in

Chapter 4 corresponds exactly to the flow of web surfers). Tomlin showed that the dual vector of the entropy optimization problem, the HOTS vector, is solution of a nonlinear fixed point equation. He then proposed a scaling algorithm to compute the HOTS vector, based on this fixed point equation.

This algorithm solves the matrix balancing problem studied among others in [Har71, EHRS85, SZ90, Sch90]. Given a $n \times n$ nonnegative matrix A , the matrix balancing problem consists in finding a matrix X of the form $X = D^{-1}AD$ with D diagonal definite positive and such that $\sum_k X_{i,k} = \sum_j X_{j,i}$ for all i . We shall compare Tomlin's HOTS algorithm with Schneider and Zenios's coordinate descent DSS algorithm [SZ90]. The main difference between these algorithms is that in coordinate descent, the scaling is done node by node in the network (i.e. in a Gauss-Seidel fashion) whereas in Tomlin's HOTS algorithm, the scaling is done all the nodes at the same time, in a Jacobi fashion.

A problem close to the matrix balancing problem is the equivalence scaling problem, where given an $m \times n$ nonnegative matrix A , we search for a matrix X of the form $X = D_1AD_2$ with D_1 and D_2 diagonal definite positive and such that X is bistochastic. The Sinkhorn-Knopp [KS67] algorithm is a famous algorithm designed for the resolution of the scaling problem. We may see HOTS algorithm as the analog of Sinkhorn-Knopp algorithm for the matrix balancing problem: both algorithms correspond to fixed point iterations on the diagonal scalings. Moreover, Smith [Smi05] and Knight [Kni08] proposed to rank web pages according to the inverse of the corresponding entry in the diagonal scaling.

However, whereas Sinkhorn-Knopp algorithm [KS67] and the coordinate descent algorithm [LT92] have been proved to converge, it does not seem that a theoretical result on the convergence of Tomlin's HOTS algorithm has been stated in previous works, although experimentations [Tom03] suggest that it is the case. Indeed, Knight [Kni08, Sec. 5] rose the fact that Tomlin did not state any convergence result for HOTS algorithm. Another algorithm for the matrix balancing problem is given in [JPS00], based on the equivalence between the matrix balancing problem and the problem of minimizing the dominant eigenvalue of an essentially nonnegative matrix under trace-preserving diagonal perturbations [JSOvdD94].

In this chapter, we prove the convergence of Tomlin's HOTS algorithm. We first study a simplified version of the algorithm that we call the ideal HOTS algorithm. It is a fixed point scaling algorithm that solves the matrix balancing problem for nonnegative irreducible matrices. We prove its convergence thanks to nonlinear Perron-Frobenius theory (Theorem 6.2). The proof methods are general and apply to a family of deformations of HOTS. Then, we address the effective HOTS algorithm, for the general case, which is the version designed by Tomlin for the ranking of web pages. Indeed the web graph is not strongly connected, which implies that the balanced matrix does not necessarily exist. The model is a nonlinear network entropy maximization problem which generalizes matrix balancing. We show in Theorem 6.5 that under mild assumptions the HOTS algorithm converges with a linear rate of convergence. The proof relies on the properties of the ideal HOTS algorithm: uniqueness of the fixed point up to an additive constant and decrease of a Lyapunov function at every step (Theorem 6.3).

We also show that Schneider and Zenios's coordinate descent algorithm can be adapted to find the ideal and effective HOTS vectors. We compare the HOTS algorithm and coordinate descent on fragments of the web graph in Section 6.5. We considered small, medium and large size problems. In all cases the respective computational costs of both algorithms were similar. As the performances of the HOTS algorithm depends on the primitivity of the adjacency matrix considered and coordinate descent does not, coordinate descent can be thought to have a wider range of applications. However, the actual implementation of the

HOTS algorithm is attractive for web scale problems: whereas coordinate descent DSS uses at each iteration (corresponding to a given web page) information from incoming and outgoing hyperlinks, the HOTS algorithm reduces to elementwise operations and left and right matrix vector products. Hence, an iteration of the HOTS algorithm can be performed without neither computing nor storing the transpose of the adjacency matrix.

We give an exact coordinate descent algorithm for the truncated scaling problem defined in [Sch89] and we extend its use to the problem of computing the HOTS vector when some bounds on the web surfers flow are known. Experimental results show that exact coordinate descent is an efficient algorithm for web scale problems and that it is faster than the inexact coordinate descent algorithm presented in [Sch90]. Finally, we remarked that the convergence rate of the effective HOTS algorithm seems to deteriorate when the size of the graph considered increases. In order to overcome this feature, we propose a normalized version of the HOTS algorithm where we maximize a relative entropy of the flow of web surfers instead of the classical entropy. A byproduct is that the associated ranking favors pages with no outlink less than Tomlin's HOTS.

The chapter is organized as follows. In Section 6.2, we prove the convergence of the ideal HOTS algorithm and we give a Lyapunov function for this algorithms. In Section 6.3, we give the convergence rate of the effective HOTS algorithm. In Section 6.4, we study the HOTS problem with bounds on the flow of web surfers. In Section 6.5, we compare various candidate algorithms to compute the HOTS vector and in Section 6.6, we give the normalized HOTS algorithm.

6.2 The ideal HOTS algorithm

The web graph is a graph constructed from the hyperlink structure of the web. Each web page is represented by a node and there is an arc between nodes i and j if and only if page i points to page j . We shall denote by A the adjacency matrix of the web graph.

There are two versions of the HOTS algorithm: an ideal version for strongly connected graphs, i.e. for irreducible adjacency matrices, and an effective version for general graphs that we will study in Section 6.3. The HOTS algorithm for irreducible matrices is designed for the resolution of the following nonlinear network flow problem. The optimization variable $\rho_{i,j}$ represents the traffic of websurfers on the hyperlink from page i to page j .

$$\begin{aligned} \max_{\rho \geq 0} & - \sum_{i,j \in [n]} \rho_{i,j} (\log(\frac{\rho_{i,j}}{A_{i,j}}) - 1) \\ & \sum_{j \in [n]} \rho_{i,j} = \sum_{j \in [n]} \rho_{j,i}, \quad \forall i \in [n] & (p_i) \\ & \sum_{i,j \in [n]} \rho_{ij} = 1 & (\mu) \end{aligned}$$

The dual problem consists in minimizing the function θ on $\mathbb{R}^n \times \mathbb{R}$ where

$$\theta(p, \mu) := \sum_{i,j \in [n]} A_{ij} e^{p_i - p_j + \mu} - \mu .$$

We use the convention that $0 \log(0) = 0$ and that $x \log(x/0) = 0$ if $x = 0$, $x \log(x/0) = +\infty$ otherwise.

If (p, μ) is a minimizer of θ , then the value of $\exp(p_i)$ is interpreted as the temperature of page i , the hotter the better. We call it the HOTS (Hyperlinked Object Temperature Scale) score.

The ideal HOTS algorithm (Algorithm 6.1) reduces to the fixed point iterations for the function f defined by

$$f(x) = \frac{1}{2}(\log(A^T e^x) - \log(Ae^{-x})) . \quad (6.1)$$

Denoting $y_i = e^{p_i}$, we can write it in multiplicative form to spare computing the exponentials and logarithms.

Algorithm 6.1 Ideal HOTS algorithm [Tom03]

Start with an initial point $y_0 \in \mathbb{R}^n$, $y_0 > 0$. Given y^k , compute y^{k+1} such that

$$y_i^{k+1} = \left(\frac{\sum_{j \in [n]} A_{j,i} y_j^k}{\sum_{l \in [n]} A_{i,l} (y_l^k)^{-1}} \right)^{\frac{1}{2}} .$$

Algorithm 6.2 Coordinate descent DSS [SZ90]

Start with an initial point $y^0 \in \mathbb{R}^n$, $y_0 > 0$. Given y^k , select a coordinate $i \in [n]$ and compute y^{k+1} such that

$$y_i^{k+1} = \left(\frac{\sum_{j \in [n]} A_{j,i} y_j^k}{\sum_{l \in [n]} A_{i,l} (y_l^k)^{-1}} \right)^{\frac{1}{2}}$$

$$y_j^{k+1} = y_j^k, \quad \forall j \neq i$$

We shall compare the HOTS algorithm with Schneider and Zenios's coordinate descent DSS algorithm (Algorithm 6.2). This is indeed a coordinate descent algorithm since for every k , we have, denoting $p_i = \log(y_i)$,

$$p_i^{k+1} = \arg \min_{x \in \mathbb{R}} \theta(p_1^k, \dots, p_{i-1}^k, x, p_{i+1}^k, \dots, p_n^k) .$$

Coordinate descent algorithms (Algorithm 6.3) are designed to solve

$$\min_{x \in \mathcal{X}} \phi(x) \quad (6.2)$$

where \mathcal{X} is a possibly unbounded box of \mathbb{R}^n and ϕ has the form $\phi(x) = \psi(Ex) + \langle b, x \rangle$, ψ is a proper closed convex function, E is an $m \times n$ matrix having no zero row and b is a vector of \mathbb{R}^n .

Proposition 6.1 ([LT92]). *Assume that the set of optimal solutions \mathcal{X}^* of (6.2) is nonempty, that the domain of ψ is open, that ψ is twice continuously differentiable on its domain and that $\nabla^2 \psi(Ex)$ is positive definite for all $x \in \mathcal{X}^*$. Let $(x^k)_k$ be a sequence generated by the coordinate descent algorithm (Algorithm 6.3), using the cyclic rule (more general rules are also possible). Then $(x^k)_k$ converges at least linearly to an element of \mathcal{X}^* .*

Algorithm 6.3 Coordinate descent

Start with an initial point $x^0 \in \mathbb{R}^n$. Given x^k , select a coordinate $i \in [n]$ and compute x^{k+1} such that

$$\begin{aligned} x_i^{k+1} &= \arg \min_{l_i \leq y \leq u_i} \phi(x_1^k, \dots, x_{i-1}^k, y, x_{i+1}^k, \dots, x_n^k) \\ x_j^{k+1} &= x_j^k, \quad \forall j \neq i \end{aligned}$$

We now study the fixed point operator f defined in (6.1).

Proposition 6.2 ([Tom03]). *A vector $p \in \mathbb{R}^n$ is a fixed point of f defined in (6.1) if and only if the couple (p, μ) with $\mu = -\log(\sum_{i,j \in [n]} A_{ij} e^{p_i - p_j})$ is a minimizer of the dual function. Moreover, in this case, denoting $D = \text{diag}(\exp(p))$, $e^\mu D A D^{-1}$ is a maximizer of the network flow problem.*

Proof. As θ is convex and differentiable, a couple (p, μ) is a minimizer if and only if it cancels the gradient. $\frac{\partial \theta}{\partial \mu}(p, \mu) = \sum_{i,j \in [n]} A_{ij} e^{p_i - p_j} e^\mu - 1$, so we have the expression of the optimal μ as a function of p . To conclude, we remark that

$$\frac{\partial \theta}{\partial p_k}(p, \mu) = \left(- \sum_{i \in [n]} A_{i,k} e^{p_i - p_k} + \sum_{j \in [n]} A_{k,j} e^{p_k - p_j} \right) e^\mu = 0$$

is equivalent to $f(p) = p$. To get back to the primal problem, we remark that the primal cost of $e^\mu D A D^{-1}$ is equal to the dual cost of (p, μ) and that it is an admissible circulation. \square

Proposition 6.3. *The map f defined in (6.1) is monotone, additively homogeneous (Definition 2.2).*

Proof. For all real λ and for all vectors p, q such that $p \leq q$, $f(\lambda + p) = \lambda + f(p)$ ($\log(e^\lambda) = \lambda$) and $f(p) \leq f(q)$ (\log and \exp are increasing functions). \square

The following result gives the conditions for the existence and uniqueness of the ideal HOTS vector.

Theorem 6.1 ([EHR85]). *There exists $v \in \mathbb{R}^n$ such that $f(v) = v$ and $\sum_{i \in [n]} v_i = 0$ if and only if A has a diagonal similarity scaling if and only if A is completely reducible.*

If in addition A is irreducible, then this vector is unique.

Corollary 6.1. *If A is completely reducible, coordinate descent DSS (Algorithm 6.2) converges linearly to a vector v such that $\text{diag}(v) A \text{diag}(v)^{-1}$ is scaled.*

To prove the convergence of the ideal HOTS algorithm (Algorithm 6.1), we use the non-linear Perron-Frobenius theory, the main theorems of which are stated in Section 2.4.

Theorem 6.2. *Let f be the map defined in (6.1). If A is irreducible and $A + A^T$ is primitive, then there exists a vector v and such that $f(v) = v$ and for all $x \in \mathbb{R}^n$,*

$$\limsup_{k \rightarrow \infty} \|f^{k+1}(x) - v\|^{1/k} \leq |\lambda_2(P)| = \max\{|\lambda|; \lambda \in \text{spectrum}(P), \lambda \neq 1\}$$

where $P = \frac{1}{2} (\text{diag}(A^T e^v)^{-1} A^T \text{diag}(e^v) + \text{diag}(A e^{-v})^{-1} A \text{diag}(e^{-v}))$. In particular, the ideal HOTS algorithm (Algorithm 6.1) converges linearly at rate $|\lambda_2(P)|$.

Proof. The iterates of the fixed point iteration defined by $p^0 = x$ and $p^{k+1} = f(p^k)$ verify $p^k = \log(y^k)$ where y^k is the k th iterate of the ideal HOTS algorithm (Algorithm 6.1) started with $y^0 = \exp(x)$. Hence, by continuous differentiability of the exponential, the rate of convergence of both versions of the algorithm is the same. By Theorem 6.1, as A is irreducible, f has a fixed point v and $\text{diag}(\exp(v))$ is solution of the matrix balancing problem associated to A . Now easy calculations show that $\nabla f = P$. As P has the same pattern as $A + A^T$, P is primitive if and only if $A + A^T$ is. The result follows from Theorem 2.3. \square

This theorem shows that the HOTS vector for the irreducible case is well defined if A is irreducible and that if $A + A^T$ is primitive, then the ideal HOTS algorithm (6.1) converges linearly to the HOTS vector.

Remark 6.1. The ideal HOTS algorithm (Algorithm 6.1) requires a primitivity assumption in order to converge that coordinate descent DSS (Algorithm 6.2) does not require. On the other hand, the convergence rate of coordinate descent DSS is not explicitly given while Theorem 6.2 gives the convergence rate of ideal HOTS.

Remark 6.2. Changing the diagonal of A does not change the optimal scaling, so we can choose a nonzero diagonal for A in the preceding theorem. This is useful when A is irreducible but not primitive.

The fixed point equation defining the ideal HOTS vector is

$$y_i = \left(\frac{\sum_j A_{j,i} y_j}{\sum_k A_{i,k} y_k^{-1}} \right)^{\frac{1}{2}}.$$

Indeed, the page i has a good HOTS score if it is linked to by pages with a good HOTS score and if it does not link to pages with a bad HOTS score.

We thus introduce the following set of fixed point ranking algorithms.

Algorithm 6.4 Deformed HOTS algorithm

Let $\alpha, \beta \geq 0$ such that $\alpha + \beta = 1$ and let $g : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ defined for all i by

$$g_i(x) = \frac{(\sum_j A_{j,i} x_j)^\alpha}{(\sum_k A_{i,k} x_k^{-1})^\beta}.$$

Given an initial point $d_0 \in \mathbb{R}^n$ and a norm $\|\cdot\|$, the deformed HOTS algorithm is defined by

$$d^{k+1} = \frac{g(d^k)}{\|g(d^k)\|}$$

Proposition 6.4. *Let $\alpha, \beta \geq 0$ such that $\alpha + \beta = 1$. If A is irreducible and $\alpha A + \beta A^T$ is primitive, then the deformed HOTS algorithm (Algorithm 6.4) converges linearly to a positive vector.*

Proof. Let $h = \log \circ g \circ \exp$. As in the proof of Theorem 6.5, the rate of convergence for the fixed point iterations with g or h is the same. The map h is monotone and additively homogeneous. For $\alpha > 0$, its graph is equal to A . Hence, for $\alpha > 0$, h has an eigenvector by

Theorem 2.1. For $\alpha = 0$, as A is irreducible, by the Perron-Frobenius theorem [BP94], A has an eigenvector x . Then $\log(x^{-1})$ is an eigenvector of h . Now

$$\nabla h(v) = \alpha \text{diag}(A^T e^v)^{-1} A^T \text{diag}(e^v) + \beta \text{diag}(Ae^{-v})^{-1} A \text{diag}(e^{-v})$$

so we have the convergence as soon as $\alpha A + \beta A^T$ is primitive by Theorem 2.3. \square

Remark 6.3. For $\alpha = \frac{1}{2}$, we have the fixed point diagonal similarity scaling, for $\alpha = 1$, we have the ranking by the Perron vector [Kee93] and for $\alpha = 0$, we have an ‘‘anti-Perron’’ score, where good pages are those that do not link to pages with a bad score.

The following result gives a global contraction factor in the case when A is positive.

Proposition 6.5. *If $k(A)$ is the contraction factor of A in Hilbert metric ($k(A) < 1$ if A is positive), then f is $\frac{k(A^T) + k(A)}{2}$ -contracting in Hilbert metric.*

Proof. Let x and y be two positive vectors such that $\eta y \leq x \leq \nu y$ elementwise. Then $\eta' A^T y \leq A^T x \leq \nu' A^T y$ with $\log(\nu'/\eta') \leq k(A^T) \log(\nu/\eta)$. We also have that $\eta'' A y^{-1} \leq A x^{-1} \leq \nu'' A y^{-1}$ with $\log((\eta'')^{-1}/(\nu'')^{-1}) \leq k(A) \log(\nu/\eta)$. Hence,

$$d(g(x), g(y)) = \log\left(\left(\frac{\nu'' \nu'}{\eta'' \eta'}\right)^{1/2}\right) \leq \frac{k(A^T) + k(A)}{2} \log\left(\frac{\nu}{\eta}\right) .$$

\square

A key technical ingredient of the convergence of the effective HOTS algorithm described in the next section will be Theorem 6.3 below showing that each iteration $p \leftarrow f(p)$ of the ideal HOTS algorithm does not increase the dual objective function.

Theorem 6.3 (Lyapunov function). $\theta(f(p)) \leq \theta(p)$

Proof. Let us denote $\psi(p, q) = \sum_{i,j} e^{p_i} A_{ij} e^{-q_j}$.

$$\begin{aligned} \psi(p, 2f(p) - p) &= \sum_{i,j} e^{p_i} A_{ij} \left(\frac{(Ae^{-p})_j}{(A^T e^p)_j} \right) e^{p_j} = \sum_j (Ae^{-p})_j e^{p_j} \\ &= \theta(p) = \psi(2f(p) - p, p) \end{aligned}$$

Now, as ψ is convex,

$$\theta(f(p)) = \psi\left(\frac{1}{2}(2f(p) - p, p) + \frac{1}{2}(p, 2f(p) - p)\right) \leq \theta(p) .$$

\square

6.3 The effective HOTS algorithm

Theorem 6.1 gives conditions for the existence and uniqueness of the HOTS vector in the ideal case. In practice the irreducibility condition does not hold for the web graph. The classical solution for this problem is to add a small positive value to the adjacency matrix [BP98, LM06] in order to get a positive matrix. Tomlin proposed an alternative approach based on the

network flow model. We consider the following nonlinear network flow problem with network given by

$$A' = \begin{bmatrix} A & 1 \\ 1^T & 0 \end{bmatrix}$$

where $\mathbf{1}$ denotes the vector with all entries equal to 1.

$$\begin{aligned} \max_{\rho \geq 0} \quad & - \sum_{i,j \in [n+1]} \rho_{i,j} (\log(\frac{\rho_{i,j}}{A'_{i,j}}) - 1) \\ & \sum_{j \in [n+1]} \rho_{i,j} = \sum_{j \in [n+1]} \rho_{j,i}, \quad \forall i \in [n+1] & (p_i) \\ & \sum_{i,j \in [n+1]} \rho_{ij} = 1 & (\mu) \\ & \sum_{j \in [n]} \rho_{n+1,j} = 1 - \alpha & (a) \\ & 1 - \alpha = \sum_{i \in [n]} \rho_{i,n+1} & (b) \end{aligned}$$

We use the conventions that $0 \log(0) = 0$ and that $x \log(x/0) = 0$ if and only if $x = 0$. In this new model, we add an artificial node connected to all the other nodes and such that the flow through this node is prescribed to be $1 - \alpha$.

The algorithm is designed for the minimization of the dual function θ where

$$\begin{aligned} \theta(p, \mu, a, b) = \sum_{i,j \in [n]} A_{ij} e^{p_i - p_j + \mu} + \sum_{i \in [n]} e^{-b - p_{n+1} + p_i + \mu} \\ + \sum_{j \in [n]} e^{a + p_{n+1} - p_j + \mu} - (1 - \alpha)a - \mu + (1 - \alpha)b. \quad (6.3) \end{aligned}$$

We first give the following counter-example, showing that the problem may be ill posed.

Counter-Example 6.1. The dual function θ may be unbounded.

Proof. Take

$$A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}.$$

We have

$$\begin{aligned} \tilde{\theta}(p) = \min_{\mu, a, b} \theta(p, \mu, a, b) &= C(\alpha) + (1 - \alpha) \log\left(\sum_{i \in [n]} e^{p_i}\right) + \\ & (1 - \alpha) \log\left(\sum_{i \in [n]} e^{-p_i}\right) + (2\alpha - 1) \log(e^{p_1 - p_2} + e^{p_2 - p_3}) \end{aligned}$$

where $C(\alpha) \in \mathbb{R}$. For all $k \in \mathbb{R}$,

$$\begin{aligned} \tilde{\theta}(-k, 0, k, 0) &= C(\alpha) + 2(1 - \alpha) \log(1 + e^k + e^{-k}) + (2\alpha - 1) \log(e^{-k} + e^{-k}) \\ &= C(\alpha) + 2(1 - \alpha)k - (2\alpha - 1)k + 2(1 - \alpha) \log(1 + e^{-k} + e^{-2k}) + (2\alpha - 1) \log(2) \end{aligned}$$

For $\alpha > \frac{3}{4}$, θ is unbounded. □

This example is indeed rather degenerate: the HOTS algorithm can only diverge because it searches the minimum of an unbounded function. Said otherwise, it tries to solve a network flow problem without any admissible flow. We shall give conditions under which there exists a HOTS vector and show that the HOTS algorithm converges to the HOTS vector when these conditions hold.

Remark 6.4. A natural idea to establish the convergence of a fixed point algorithm is to show that it is a contraction in Hilbert metric. In the case of the effective HOTS algorithm, even when the matrix A is positive, the fixed point algorithm may not be a contraction (take a perturbation of Counter-example 6.1).

Lemma 6.1 ([Tom03]). *For any $p \in \mathbb{R}^{n+1}$, the minimum of $\theta(p, \mu, a, b)$ with respect to μ , a and b is unique and given by*

$$\begin{aligned}\mu &= \log\left(\frac{2\alpha - 1}{\sum_{i,j \in [n]} A_{i,j} e^{p_i - p_j}}\right), \\ a &= \log\left(\frac{1 - \alpha}{2\alpha - 1} \frac{\sum_{i,j \in [n]} A_{i,j} e^{p_i - p_j}}{\sum_{j \in [n]} e^{p_{n+1} - p_j}}\right), \\ b &= -\log\left(\frac{1 - \alpha}{2\alpha - 1} \frac{\sum_{i,j \in [n]} A_{i,j} e^{p_i - p_j}}{\sum_{i \in [n]} e^{p_i - p_{n+1}}}\right).\end{aligned}$$

Proof. The function $\theta(p, \cdot, \cdot, \cdot)$ is convex and differentiable so the optimality condition is just that the gradient is zero. One can easily see that the only triple that cancels the gradient is the one given in the lemma. \square

We denote $\lambda = (\mu, a, b)$ and $\lambda(p)$ the solution of the minimization of $\theta(p, \lambda)$ with respect to λ . For $\lambda \in \mathbb{R}^3$, we denote

$$f_i^\lambda(p) = \frac{1}{2} \left(\log\left(\sum_{j \in [n]} A_{j,i} e^{p_j} + e^{p_{n+1} + a}\right) - \log\left(\sum_{k \in [n]} A_{i,k} e^{-p_k} + e^{-p_{n+1} - b}\right) \right). \quad (6.4)$$

We also define $g^\lambda = \exp \circ f^\lambda \circ \log$:

$$g_i^\lambda(y) = \left(\frac{\sum_{j \in [n]} A_{j,i} y_j + e^a y_{n+1}}{\sum_{k \in [n]} A_{i,k} (y_k)^{-1} + e^{-b} (y_{n+1})^{-1}} \right)^{\frac{1}{2}}.$$

Algorithm 6.5 Effective HOTS algorithm

Given an initial point $y^0 \in \mathbb{R}^n$, the effective HOTS algorithm is defined by

$$y^{k+1} = g^{\lambda(\log(y^k))}(y^k)$$

Proposition 6.6. *Let us denote $\tilde{\theta}(p) = \min_{p_{n+1}, \mu, a, b} \theta((p, p_{n+1}), \mu, a, b)$ and for all $l \in [n]$,*

$$d_l = \frac{e^{p_l} (\sum_{j \in [n]} e^{-p_j}) (\sum_{i,j \in [n]} A_{i,j} e^{p_i - p_j})}{(2\alpha - 1) (\sum_{i \in [n]} A_{il} e^{p_i}) (\sum_{j \in [n]} e^{-p_j}) + (1 - \alpha) \sum_{i,j \in [n]} A_{i,j} e^{p_i - p_j}} > 0. \quad (6.5)$$

Then

$$\tilde{\theta}(p) = C(\alpha) + \phi(-p) + \phi(p) + (2\alpha - 1) \log\left(\sum_{i,j \in [n]} A_{i,j} e^{p_i - p_j}\right)$$

where $C(\alpha) = 1 - 2(1 - \alpha) \log(1 - \alpha) - (2\alpha - 1) \log(2\alpha - 1)$ and $\phi(p) = (1 - \alpha) \log(\sum_{i \in [n]} e^{p_i})$. Moreover, for all $l \in [n]$, we have

$$f_l^{\lambda(p)}(p) = p_l - \frac{1}{2} \log\left(1 + d_l \frac{\partial \tilde{\theta}}{\partial p_l}(p)\right) .$$

Proof. From the expressions of a_{n+1} , b_{n+1} and μ at the optimum, respectively given by $e^{a_{n+1}} = \frac{1-\alpha}{\sum_{j \in [n]} e^{-p_j}} e^{-\mu}$, $e^{-b_{n+1}} = \frac{1-\alpha}{\sum_{i \in [n]} e^{p_i}} e^{-\mu}$ and $e^\mu = \frac{2\alpha-1}{\sum_{i,j \in [n]} A_{ij} e^{p_i - p_j}}$, we can write $\tilde{\theta}$ as a function of p only:

$$\tilde{\theta}(p) = \min_{a_{n+1}, b_{n+1}, \mu} \theta(p, \mu, a_{n+1}, b_{n+1}) = C(\alpha) + \phi(-p) + \phi(p) + (2\alpha - 1) \log\left(\sum_{i,j \in [n]} A_{i,j} e^{p_i - p_j}\right)$$

where $C(\alpha)$ and ϕ are given in the proposition. Its gradient is given by

$$\frac{\partial \tilde{\theta}}{\partial p_l}(p) = -(1 - \alpha) \frac{e^{-p_l}}{\sum_j e^{-p_j}} + (1 - \alpha) \frac{e^{p_l}}{\sum_i e^{p_i}} - (2\alpha - 1) \frac{\sum_i A_{il} e^{p_i - p_l}}{\sum_{i,j} A_{ij} e^{p_i - p_j}} + (2\alpha - 1) \frac{\sum_j A_{lj} e^{p_l - p_j}}{\sum_{i,j} A_{ij} e^{p_i - p_j}}$$

This equality can be also written as

$$e^{2p_l} \left((2\alpha - 1) \frac{\sum_j A_{lj} e^{-p_j}}{\sum_{i,j} A_{ij} e^{p_i - p_j}} + \frac{1 - \alpha}{\sum_i e^{p_i}} \right) = (2\alpha - 1) \frac{\sum_i A_{il} e^{p_i}}{\sum_{i,j} A_{ij} e^{p_i - p_j}} + \frac{1 - \alpha}{\sum_j e^{-p_j}} + e^{p_l} \frac{\partial \tilde{\theta}}{\partial p_l}(p)$$

which yields for all p in \mathbb{R}^n ,

$$p_l = \frac{1}{2} \log \left(\sum_i A_{il} e^{p_i} + \frac{\sum_{i,j} A_{ij} e^{p_i - p_j}}{2\alpha - 1} \left(\frac{1 - \alpha}{\sum_j e^{-p_j}} + e^{p_l} \frac{\partial \tilde{\theta}}{\partial p_l}(p) \right) \right) - \frac{1}{2} \log \left(\left(\sum_j A_{lj} e^{-p_j} \right) + \frac{1 - \alpha}{2\alpha - 1} \frac{\sum_{i,j} A_{ij} e^{p_i - p_j}}{\sum_i e^{p_i}} \right) .$$

Now,

$$f_l^{\lambda(p)}(p) = \frac{1}{2} \log \left(\sum_i A_{il} e^{p_i} + \frac{1 - \alpha}{2\alpha - 1} \frac{\sum_{i,j} A_{ij} e^{p_i - p_j}}{\sum_j e^{p_j}} \right) - \frac{1}{2} \log \left(\sum_j A_{lj} e^{-p_j} + \frac{1 - \alpha}{2\alpha - 1} \frac{\sum_{i,j} A_{ij} e^{p_i - p_j}}{\sum_i e^{p_i}} \right) .$$

Using the formula $\log(A) = \log(A + B) - \log(1 + B/A)$, we may also write it as

$$f_l^{\lambda(p)}(p) = p_l - \frac{1}{2} \log\left(1 + d_l \frac{\partial \tilde{\theta}}{\partial p_l}(p)\right) . \quad \square$$

We can see that the equation $f^{\lambda(p)}(p) = p$ is equivalent to $\frac{\partial \tilde{\theta}}{\partial p_l}(p) = 0$ but also that successively applying the function $(p \mapsto f^{\lambda(p)}(p))$ corresponds to a descent algorithm.

The next proposition gives information on the spectrum of the Hessian of the function θ .

Proposition 6.7. *The Hessian of the function $\tilde{\theta} : p \mapsto \tilde{\theta}(p) = \min_{p_{n+1}, \mu, a, b} \theta((p, p_{n+1}), \mu, a, b)$ is symmetric semi-definite with spectral norm smaller than 4. Its nullspace has dimension 1 exactly for all p and a basis of this nullspace is given by the vector e , with all entries equal to 1.*

Proof. As θ is convex, its Hessian matrix is clearly symmetric semi-definite.

Now, let $\phi : x \mapsto \log(\sum_i e^{x_i})$ be the log-sum-exp function.

We have $y^T \nabla^2 \phi(x) y = \frac{\sum_i y_i^2 e^{x_i}}{\sum_k e^{x_k}} - \frac{(\sum_i y_i e^{x_i})^2}{(\sum_k e^{x_k})^2}$. This expression is strictly positive for any non constant y , because a constant y is the only equality case of the Cauchy Schwartz inequality $\sum_i y_i e^{x_i/2} e^{x_i/2} \leq (\sum_i e^{x_k})^{1/2} (\sum_i y_i^2 e^{x_i})^{1/2}$. As the function θ is the sum of ϕ (the third term of (6.6)) and of convex functions, it inherits the strict convexity property on spaces without constants. This development even shows that the kernel of $\nabla^2 \theta(p)$ is of dimension at most 1 for all p . Finally, as θ is invariant by addition of a constant, the vector e is clearly part of the nullspace of its Hessian.

For the norm of the Hessian matrix, we introduce the linear function $Z : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ such that $(Zp)_{i,j} = p_i - p_j$ and $\tilde{\phi} : z \mapsto \log(\sum_{k \in [n] \times [n]} A_k e^{z_k})$. Then $\theta(p) = C(\alpha) + (1 - \alpha)\phi(p) + (1 - \alpha)\phi(-p) + (2\alpha - 1)\tilde{\phi}(Zp)$.

$$0 \leq y^T \nabla^2 \phi(x) y = \frac{\sum_i y_i^2 e^{x_i}}{\sum_k e^{x_k}} - \frac{(\sum_i y_i e^{x_i})^2}{(\sum_k e^{x_k})^2} \leq \frac{\sum_i y_i^2 e^{x_i}}{\sum_k e^{x_k}} \leq \|y\|_\infty^2 \leq \|y\|_2^2. \text{ Thus } \|\nabla^2 \phi\|_2 \leq 1.$$

By similar calculations, one gets $y^T Z^T \nabla^2 \tilde{\phi}(x) Z y \leq \|Zy\|_\infty^2$. As $\|Zy\|_\infty = \max_{i,j} |y_i - y_j| \leq 2\|y\|_\infty \leq 2\|y\|_2$, we have that $\|\nabla^2 \theta\|_2 \leq (1 - \alpha) + (1 - \alpha) + (2\alpha - 1) \times 4 = 6\alpha - 2 < 4$. Finally, for symmetric matrices, the spectral norm and the operator 2-norm are equal. \square

Lemma 6.2. *Let $d \in \mathbb{R}^n$ such that for all i , $d_i > 0$. The reduced function $\tilde{\theta}$ defined by $\tilde{\theta}(p) = \min_{p_{n+1}, \mu, a, b} \theta((p, p_{n+1}), \mu, a, b)$ is strictly convex on $H = \{x \in \mathbb{R}^n \mid \sum_{i \in [n]} d_i x_i = 0\}$. In particular, there exists at most one HOTS vector up to an additive constant.*

Proof. Like in the proof of Proposition 6.7, as H does not contain the constant vectors, $\tilde{\theta}$ is strictly convex on H . We conclude that the minimum of $\tilde{\theta}$ on H is unique if it exists. We can then extend this result to the whole space since $\theta(\eta + p) = \theta(p)$ for all real number η . \square

Lemma 6.3. *Let d and $\tilde{\theta}$ be as in Lemma 6.2. The function $\tilde{\theta}$ is coercive on the hyperplane $\{x \in \mathbb{R}^{n+1} \mid \sum_{i \in [n+1]} d_i x_i = 0\}$ if and only if there exists a primal solution with the same pattern as A' .*

Proof. If the function θ is coercive, there exists a dual solution and thus there also exists a primal solution with the same pattern as A' .

If there exists a primal solution with the same pattern as A' , the constraint qualification conditions are satisfied [Sch89], and there exists a dual solution. By Lemma 6.2, $\tilde{\theta}$ is strictly convex on the hyperplane $\{x \in \mathbb{R}^{n+1} \mid \sum_{i \in [n+1]} d_i x_i = 0\}$. Thus it is necessarily coercive on this hyperplane. \square

Lemma 6.4. *If $A \neq 0$, then for any fixed λ , the iterative algorithm consisting in successive applications of f^λ , defined in (6.4), converges to a minimizer of the function $(p \mapsto \theta(p, \lambda))$. Moreover, this minimizer is unique up to an additive constant.*

Proof. The map f^λ corresponds to the ideal HOTS fixed point operator (6.1) for the matrix

$$\begin{bmatrix} A & e^a \mathbf{1} \\ e^b \mathbf{1}^T & 0 \end{bmatrix}$$

where $\mathbf{1}$ denotes the vector with all entries equal to 1. As this matrix is primitive as soon as $A \neq 0$, Theorem 6.2 and Proposition 6.2 apply. \square

Theorem 6.4. *Let F defined by $F(p) = f^{\lambda(p)}(p)$ as in (6.4). Let p^* be the logarithm of a HOTS vector defined by $p^* = F(p^*)$. The matrix $\nabla F(p^*)$ has all its eigenvalues in the real interval $(-1, 1]$ and the eigenvalue 1 is simple.*

Proof. Let us denote $\gamma = \frac{1-\alpha}{2\alpha-1}$, $e^a = \gamma \frac{\sum_{i,j \in [n]} A_{ij} e^{p_i - p_j}}{\sum_j e^{-p_j}}$, $e^{-b} = \gamma \frac{\sum_{i,j \in [n]} A_{ij} e^{p_i - p_j}}{\sum_i e^{p_i}}$. Then for all $k \in [n]$,

$$F_k(p) = \frac{1}{2} \log\left(\sum_{i \in [n]} A_{ik} e^{p_i} + e^a\right) - \frac{1}{2} \log\left(\sum_{j \in [n]} A_{kj} e^{-p_j} + e^{-b}\right)$$

$$F_{n+1}(p) = \frac{1}{2} \log\left(\sum_{i \in [n]} e^{p_i}\right) - \frac{1}{2} \log\left(\sum_{j \in [n]} e^{-p_j}\right).$$

As no coordinate of $F(p)$ depends on p_{n+1} , we may consider the reduced function that we shall still denote F and such that to $p \in \mathbb{R}^n$ associates $F(p, 0)$. The eigenvalues of the original function are 0 and the eigenvalues of the reduced function.

First, by Proposition 6.6, we have

$$F(p) = p - \frac{1}{2} \log\left(1 + \text{diag}(d) \frac{\partial \tilde{\theta}}{\partial p}\right)$$

where $\tilde{\theta}(p) = \min_{p_{n+1}, \mu, a, b} \theta((p, p_{n+1}), \mu, a, b)$ and $d_l = \frac{1}{2\alpha-1} \frac{e^{p_l} (\sum_{i,j \in [n]} A_{ij} e^{p_i - p_j})}{\sum_{i \in [n]} A_{il} e^{p_i} + e^a} > 0$. Differentiating this equality, we deduce that $\nabla F = I_n - \frac{1}{2} \text{diag}(d) \nabla^2 \tilde{\theta}$. Let λ be an eigenvalue of ∇F . This means that there exist a vector x such that

$$\lambda x = \nabla F x = x - \frac{1}{2} \text{diag}(d) \nabla^2 \tilde{\theta} x$$

$$\nabla^2 \tilde{\theta} x = 2(1 - \lambda) \text{diag}(d^{-1}) x$$

This is a generalized eigenvalue problem with $\nabla^2 \tilde{\theta}$ symmetric semi-definite positive by convexity of $\tilde{\theta}$ and $\text{diag}(d^{-1})$ diagonal definite positive. Hence $2(1 - \lambda)$ is necessarily a nonnegative real number and λ is real and smaller than 1. Also, if $\lambda = 1$, this means that x is the vector with all its entries equal to 1 (by Proposition 6.7) and thus λ is simple.

We shall now show that all the eigenvalues of $\nabla F(p^*)$ are strictly greater than -1 . Differentiating the expression of a , we get

$$\frac{\partial e^a}{\partial p_l}(p) = \gamma \left(-\frac{\sum_{i \in [n]} A_{il} e^{p_i - p_l}}{\sum_{j \in [n]} e^{-p_j}} + \frac{\sum_{j \in [n]} A_{lj} e^{p_l - p_j}}{\sum_{j \in [n]} e^{-p_j}} + \frac{\sum_{i,j \in [n]} A_{ij} e^{p_i - p_j}}{(\sum_{j \in [n]} e^{-p_j})^2} e^{-p_l} \right)$$

But as p^* is a fixed point of F , it satisfies the equality

$$\sum_{i \in [n]} A_{il} e^{p_i^* - p_l^*} + e^{a^* - p_l^*} = \sum_{j \in [n]} A_{lj} e^{p_l^* - p_j^*} + e^{-b^* + p_l^*}$$

which can be rewritten as

$$\sum_{j \in [n]} A_{lj} e^{p_l^* - p_j^*} - \sum_{i \in [n]} A_{il} e^{p_i^* - p_l^*} = \gamma \sum_{i, j \in [n]} A_{ij} e^{p_i^* - p_j^*} \left(\frac{e^{-p_l^*}}{\sum_{j \in [n]} e^{-p_j^*}} - \frac{e^{p_l^*}}{\sum_{i \in [n]} e^{p_i^*}} \right).$$

Hence

$$\frac{\partial e^a}{\partial p_l}(p^*) = \gamma \sum_{i, j \in [n]} A_{ij} e^{p_i^* - p_j^*} \left(\gamma \frac{e^{-p_l^*}}{(\sum_j e^{-p_j^*})^2} - \gamma \frac{e^{p_l^*}}{\sum_i e^{p_i^*} \sum_j e^{-p_j^*}} + \frac{e^{-p_l^*}}{(\sum_j e^{-p_j^*})^2} \right)$$

Let us introduce d' such that

$$d'_k{}^{-1} = \sum_{i \in [n]} A_{ik} e^{p_i^* - p_k^*} + e^{a^* - p_k^*} = \sum_{j \in [n]} A_{kj} e^{p_k^* - p_j^*} + e^{-b^* + p_k^*}. \quad (6.6)$$

Doing the same for e^{-b} as for e^a and differentiating F , we get

$$\begin{aligned} \frac{\partial F_k}{\partial p_l}(p^*) &= \frac{1}{2} d'_k A_{l,k} e^{p_l^* - p_k^*} + \frac{1}{2} d'_k A_{k,l} e^{p_k^* - p_l^*} + \frac{1}{2} \gamma \sum_{i, j \in [n]} A_{ij} e^{p_i^* - p_j^*} d'_k \times \\ &\quad \left((1 + \gamma) \frac{e^{-p_l^* - p_k^*}}{(\sum_j e^{-p_j^*})^2} - \gamma \frac{e^{p_l^* - p_k^*}}{\sum_i e^{p_i^*} \sum_j e^{-p_j^*}} - \gamma \frac{e^{p_k^* - p_l^*}}{\sum_i e^{p_i^*} \sum_j e^{-p_j^*}} + (1 + \gamma) \frac{e^{p_l^* + p_k^*}}{(\sum_i e^{p_i^*})^2} \right). \end{aligned}$$

We can now decompose $\frac{\partial F}{\partial p}$ as

$$\frac{\partial F}{\partial p} = D'S + D'R$$

where $D' = \text{diag}(d')$, S is a symmetric matrix with nonnegative entries

$$S_{k,l} = \frac{1}{2} A_{k,l} e^{p_l - p_k} + \frac{1}{2} A_{l,k} e^{p_k - p_l} + \frac{1}{2} \gamma \sum_{i, j} A_{ij} e^{p_i - p_j} \left(\frac{e^{-p_l - p_k}}{(\sum_j e^{-p_j})^2} + \frac{e^{p_l + p_k}}{(\sum_i e^{p_i})^2} \right)$$

and R is the following symmetric rank 1 matrix

$$R_{k,l} = \frac{1}{2} \gamma^2 \sum_{i, j} A_{ij} e^{p_i - p_j} \left(\frac{e^{-p_k}}{\sum_j e^{-p_j}} - \frac{e^{p_k}}{\sum_i e^{p_i}} \right) \left(\frac{e^{-p_l}}{\sum_j e^{-p_j}} - \frac{e^{p_l}}{\sum_i e^{p_i}} \right)$$

The nonnegative matrix $D'S$ verifies that for all k , $\sum_l d'_k S_{k,l} = 1$, thus by the Perron-Frobenius theorem [BP94], we have exhibited a Perron vector and the spectral radius of the matrix is 1. Moreover, $D'S$ is positive, so every other of its eigenvalues has a modulus strictly smaller than 1.

The matrix $D'R$ is a rank 1 matrix and its only nonzero eigenvalue is positive. Indeed it is equal to $\frac{1}{2} \gamma^2 \sum_{i, j} A_{ij} e^{p_i - p_j} \sum_k \left(\frac{e^{-p_k}}{\sum_j e^{-p_j}} - \frac{e^{p_k}}{\sum_i e^{p_i}} \right)^2 d'_k$.

Let λ be an eigenvalue of $\frac{\partial F}{\partial p} = D'S + D'R$. It then also an eigenvalue of the similar matrix $(D')^{1/2} S (D')^{1/2} + (D')^{1/2} R (D')^{1/2}$, which is symmetric. Hence,

$$\lambda \geq \min_{x \in \mathbb{R}^n: \|x\|_2=1} x^T (D')^{1/2} S (D')^{1/2} x + x^T (D')^{1/2} R (D')^{1/2} x$$

As the spectral radius of $D'S$ is 1, the same is true for $(D')^{1/2}S(D')^{1/2}$ and for all vector x , $x^T(D')^{1/2}S(D')^{1/2}x > -\|x\|_2^2$. As $D'R$ has only nonnegative eigenvalues, $(D')^{1/2}R(D')^{1/2}$ is semi-definite positive and $x^T(D')^{1/2}R(D')^{1/2}x \geq 0$ for all x . As a conclusion,

$$\lambda \geq \min_{x \in \mathbb{R}^n: \|x\|_2=1} x^T(D')^{1/2}S(D')^{1/2}x + x^T(D')^{1/2}R(D')^{1/2}x > -1$$

□

Theorem 6.5. *Let $F(p) = f^{\lambda(p)}(p)$ as in (6.4). If there exists a primal feasible point with the same pattern as A' , then the effective HOTS algorithm (Algorithm 6.5) converges to a HOTS vector e^{p^*} (unique up to a multiplicative constant) linearly at a rate $|\lambda_2(\nabla F(p^*))| = \max\{|\lambda|; \lambda \in \text{spectrum}(\nabla F(p^*)), \lambda \neq 1\}$,*

Proof. Let \tilde{F} be the map defined by $\tilde{F}(p) = F(p) - \frac{1}{1^T d'^{-1}} \sum_{i \in [n+1]} (d'_i)^{-1} F_i(p)$, with d'_i defined in (6.6) in the proof of Theorem 6.4 for $i \in [n]$ and $(d'_{n+1})^{-1} = 0$. For all k , let p_k be the k th iterate of the HOTS algorithm, i.e. $p_{k+1} = F(p_k)$, and let $\lambda_k = \lambda(p_k)$. We also define q_k by $q_0 = p_0$ and $q_{k+1} = \tilde{F}(q_k)$. By Theorem 6.3 and by definition of λ_{k+1} , we have $\theta(p_k, \lambda_k) \geq \theta(p_{k+1}, \lambda_k) \geq \theta(p_{k+1}, \lambda_{k+1})$. As $q_k - p_k$ is proportional to the vector with all entries equal to 1, $\lambda(p_k) = \lambda(q_k) = \lambda_k$ and $\theta(q_k, \lambda) = \theta(p_k, \lambda)$ for all k and λ . Hence

$$\theta(q_k, \lambda_k) \geq \theta(q_{k+1}, \lambda_k) \geq \theta(q_{k+1}, \lambda_{k+1}) . \quad (6.7)$$

Now, for all k , $q_k \in H = \{x \in \mathbb{R}^{n+1} \mid \sum_{i \in [n+1]} (d'_i)^{-1} x_i = 0\}$. As by Lemma 6.3, $\tilde{\theta}$ is coercive on H , θ is bounded from below and $(\theta(q_k, \lambda_k))_k$ converges to, say, $\bar{\theta}$. Moreover, the sequence $(q_k)_k$ must be bounded. By continuity of the function $\lambda(\cdot)$, $(\lambda_k)_k$ is also bounded. Hence, they have limit points.

Let \bar{q} be a limit point of $(q_k)_k$ and $\bar{\lambda} = \lambda(\bar{q})$. For all $\epsilon > 0$ and $K > 0$, there exists $k \geq K$ and $k' \geq k + 1$ such that $\|q_k - \bar{q}\| \leq \epsilon$, $\|q_{k'} - \bar{q}\| \leq \epsilon$. By (6.7),

$$\theta(q_k, \lambda_k) \geq \theta(q_{k+1}, \lambda_k) \geq \theta(q_{k+1}, \lambda(q_{k+1})) \geq \theta(q_{k'}, \lambda_{k'})$$

where $q_{k+1} = \tilde{F}(q_k)$. When ϵ tends to 0 and K tends to infinity, we get with $\bar{q} = \tilde{F}(\bar{q})$,

$$\theta(\bar{q}, \bar{\lambda}) \geq \theta(\bar{q}, \bar{\lambda}) \geq \theta(\bar{q}, \lambda(\bar{q})) \geq \theta(\bar{q}, \bar{\lambda}) .$$

In particular, $\theta(\bar{q}, \bar{\lambda}) = \theta(\bar{q}, \lambda(\bar{q}))$. This implies by Lemma 6.1 that $\bar{\lambda} \in \arg \min_{\lambda} \theta(\bar{q}, \lambda)$ and thus $\bar{\lambda} = \lambda(\bar{q})$ by uniqueness of the minimizer.

Similarly, \bar{q} is also a limit point of $(q_k)_k$ and we may consider the sequence $(u_k)_k$ such that $u_k = (\tilde{F})^k(\bar{q}) = (f^{\bar{\lambda}})^k(\bar{q}) - \frac{1}{1^T d'^{-1}} \sum_{i \in [n+1]} d'_i{}^{-1} ((f^{\bar{\lambda}})^k(\bar{q}))_i$. Iterating the argument of the preceding paragraph, for any k , $\lambda(u_k) = \bar{\lambda}$ and u_k is a limit point of $(q_k)_k$. Now, by Lemma 6.4, the sequence (u_k) converges to $q^* \in \arg \min_q \theta(q, \bar{\lambda})$. As we also have $\lambda(q^*) = \bar{\lambda}$, we conclude that $(q^*, \bar{\lambda})$ is a minimizer of θ and that there exists a limit point of $(q_k, \lambda_k)_k$ that minimizes θ . Now, as $(\theta(q_k, \lambda_k))_k$ is decreasing, all the limit points of (q_k) minimize θ . The uniqueness of the minimizer of θ on H (Lemma 6.2) gives the convergence of the effective HOTS algorithm to the HOTS vector in the projective space, that is the convergence of the sequence $(q_k)_k$.

We shall now prove that the sequence $(p_k)_{k \geq 0}$ indeed converges. Let us denote by $\rho = \max\{|\lambda|; \lambda \in \text{spectrum}(\nabla F(q^*)), \lambda \neq 1\}$. By Theorem 6.4, we know $\rho < 1$. Denoting $(\lambda_i, u_i, v_i)_{i \in [n+1]}$, the eigenvalues and eigenvectors of $\nabla F(q^*) = \sum_{i=1}^{n+1} \lambda_i u_i v_i^T$, we

have $\nabla \tilde{F}(q^*) = \sum_{i=2}^{n+1} \lambda_i u_i v_i^T + \frac{1(d'^{-1})^T}{1^T d'^{-1}} - \frac{1}{1^T d'^{-1}} (d'^{-1})^T \frac{1(d'^{-1})^T}{1^T d'^{-1}} = \sum_{i=2}^{n+1} \lambda_i u_i v_i^T$. Hence by [Ost55], for all $\epsilon' > 0$, there exists a norm $\|\cdot\|$ such that for all $x \in \mathbb{R}^{n+1}$, $\|\nabla \tilde{F}(q^*)x\| = \|\sum_{i=2}^n \lambda_i u_i v_i^T x\| \leq (\rho + \epsilon')\|x\|$.

Thus for $x \in \mathbb{R}^{n+1}$ sufficiently close to q^* ,

$$\|\tilde{F}(x) - q^*\| = \|\tilde{F}(x) - \tilde{F}(q^*)\| \leq (1 + \epsilon'/2)\|\nabla F(q^*)(x - q^*)\| \leq (\rho + \epsilon')\|x - q^*\| .$$

We deduce that (q_k) converges linearly at rate ρ to q^* . Now for all k , we have

$$\begin{aligned} p_k &= q_k + \sum_{l=0}^{k-1} \frac{1}{1^T d'^{-1}} \sum_{i \in [n+1]} d'_i{}^{-1} F_i(q_l) \\ &= q_k + \sum_{l=0}^{k-1} \frac{1}{1^T d'^{-1}} \sum_{i \in [n+1]} d'_i{}^{-1} F_i(q_l) - F_i(q^*) = q_k + \sum_{l=0}^{k-1} \eta_l , \end{aligned}$$

where $|\eta_l| = O(\|F(q_l) - F(q^*)\|) = O(\|q_l - q^*\|) = O(\rho^l)$. Hence $\sum_{l=0}^{k-1} \eta_l$ is summable and converges linearly at rate ρ . Finally, (p_k) converges linearly at rate ρ . Like in the proof of Theorem 6.5 we deduce the convergence of the sequence $(\exp(p_k))$ linearly at rate ρ to a HOTS vector. \square

Remark 6.5. If instead of considering the following nonlinear network flow problem where the network is given by

$$A' = \begin{bmatrix} A & 1 \\ 1^T & 0 \end{bmatrix}$$

we take

$$A'' = \begin{bmatrix} A & 1 \\ 1^T & 1 \end{bmatrix}$$

then there always exists a primal feasible point with the same pattern as A'' and the modified HOTS algorithm always converges (even for $A = 0$). All the proofs carry over, with only a minor change in the expression of $\lambda(p)$.

The last result shows that coordinate descent is an alternative algorithm for the computation of the effective HOTS vector.

Proposition 6.8. *If there exists a primal feasible point with the same pattern as A' , the coordinate descent algorithm applied to the unrestricted minimization of the dual function θ defined in (6.3) and choosing coordinates in a cyclic order converges linearly to a HOTS vector.*

Proof. If there exists a primal feasible point with the same pattern as A' , then the set of minimizers of θ is nonempty by Lemmas 6.2 and 6.3. The function θ has the required form with $\psi(x) = \sum_{i,j} A_{i,j} \exp(x_{i,j})$. The Hessian of ψ is clearly definite positive for all x . Thus the hypotheses of Proposition 6.1 are verified and the result follows. \square

6.4 An exact coordinate descent for the truncated scaling problem

Truncated scaling problems were introduced by Schneider in [Sch89, Sch90] in order to generalize both matrix balancing and row-column equivalence scaling. Given an $n \times n$ matrix A and bounds L and U such that $L_{i,j} \leq U_{i,j}$, the truncated scaling problem consists in finding a matrix X of the form $X = T(D^{-1}AD)$ with D diagonal definite positive, T the truncation operator $T_{i,j}(X) = \max(\min(U_{i,j}, X_{i,j}), L_{i,j})$ and such that $\sum_k X_{i,k} = \sum_j X_{j,i}$ for all i . This problem is equivalent to the following optimization problem

$$\begin{aligned} \max_{\rho \geq 0} \quad & -\sum_{i,j \in [n]} \rho_{i,j} (\log(\frac{\rho_{i,j}}{A_{i,j}}) - 1) \\ \sum_{j \in [n]} \rho_{i,j} = \sum_{j \in [n]} \rho_{j,i}, \quad & \forall i \in [n] \quad (p_i) \\ L_{i,j} \leq \rho_{i,j} \leq U_{i,j}, \quad & \forall i, j \in [n] \quad (\eta_{i,j}, \zeta_{i,j}) \end{aligned}$$

If the bounds satisfy $L_{i,j} = 0$ and $U_{i,j} = +\infty$ for all i, j , then we have a matrix balancing problem. The reduction of row-column equivalence scaling to truncated scaling lies in a graph transformation described by Schneider in [Sch89], Lemma 1.

The dual function can take two forms. In [Sch89], Schneider proposes to relax the equality constraints and to let the bound constraints in the objective function. One gets a dual function of the form

$$\Psi(p) = \sum_{i,j \in [n]} \psi_{i,j}^*(p_i - p_j) \quad (6.8)$$

where each $\psi_{i,j}^* : \mathbb{R} \rightarrow \mathbb{R}$ is convex. Then one can perform an inexact coordinate descent where at each step the minimization along the coordinate is not necessarily exact.

Here, we shall study the choice of relaxing all the constraints. This approach has been proposed in [ZC91] for another generalization of the row-column equivalence scaling problem (but this generalization does not include truncated scaling). Then, we get the following dual function:

$$\theta(p, \eta, \zeta) = \sum_{i,j \in [n]} \phi_{i,j}^*(p_i - p_j + \eta_{i,j} - \zeta_{i,j}) - \sum_{i,j \in [n]} L_{i,j} \eta_{i,j} + \sum_{i,j \in [n]} U_{i,j} \zeta_{i,j}, \quad (6.9)$$

where $\phi_{i,j}^*(t) = A_{i,j} e^t$. We shall minimize θ with unrestricted p and nonnegative η and ζ . As in [CZ91, ZC91], we shall show in Proposition 6.9 that exact expressions of the minimizers along one single coordinate exist.

Proposition 6.9. *Given $p \in \mathbb{R}^n$, the minimizers of $\min_{\eta \geq 0, \zeta \geq 0} \theta(p, \eta, \zeta)$ are given for all i and j by*

$$\begin{aligned} \exp(\eta_{i,j}) &= \max\left(\frac{L_{i,j}}{A_{i,j} e^{p_i - p_j}}, 1\right) \\ \exp(-\zeta_{i,j}) &= \min\left(\frac{U_{i,j}}{A_{i,j} e^{p_i - p_j}}, 1\right) \end{aligned}$$

Proof. The proofs for η and ζ are symmetric, so we only do the one for η .

$$\frac{\partial \theta}{\partial \eta_{i,j}} = A_{i,j} e^{p_i - p_j + \eta_{i,j} - \zeta_{i,j}} - L_{i,j}$$

Two cases may occur: either $\frac{\partial \theta}{\partial \eta_{i,j}} = 0$ and $\eta_{i,j} \geq 0$ or $\eta_{i,j} = 0$ and $\frac{\partial \theta}{\partial \eta_{i,j}} \geq 0$. Hence, either $\exp(\eta_{i,j}) = \frac{L_{i,j}}{A_{i,j}e^{p_i - p_j - \zeta_{i,j}}}$ and $\exp(\eta_{i,j}) \geq 1$ or we have $\exp(\eta_{i,j}) = 1$ and $1 \geq \frac{L_{i,j}}{A_{i,j}e^{p_i - p_j - \zeta_{i,j}}} \geq \frac{L_{i,j}}{A_{i,j}e^{p_i - p_j}}$. We thus have the result if $\eta_{i,j}$ and $\zeta_{i,j}$ are not positive together.

Now suppose that $\eta_{i,j} > 0$ and $\zeta_{i,j} > 0$. In this case, we have $\exp(\eta_{i,j}) = \frac{L_{i,j}}{A_{i,j}e^{p_i - p_j - \zeta_{i,j}}}$ and $\exp(-\zeta_{i,j}) = \frac{U_{i,j}}{A_{i,j}e^{p_i - p_j + \eta_{i,j}}}$. This implies that $U_{i,j} = L_{i,j}$. Thus the two bound constraints are in fact an equality constraint and we shall consider the unconstrained multiplier $\xi_{i,j} = \eta_{i,j} - \zeta_{i,j}$ instead of the two former multipliers. Then $\xi_{i,j}$ verifies $\exp(\xi_{i,j}) = \frac{L_{i,j}}{A_{i,j}e^{p_i - p_j}}$. It is unique and can be decomposed into $\zeta_{i,j}$ and $\eta_{i,j}$ as in the proposition. \square

This proposition shows that we do not need to store the values of $\eta_{i,j}$ and $\zeta_{i,j}$. Indeed, for fixed p and λ , if $(\eta, \zeta) = \arg \min_{\eta' \geq 0, \zeta' \geq 0} \theta(p, \lambda, \eta', \zeta')$, then if we denote by $\text{mid}(x, a, b) = \max(\min(x, a), b)$

$$\begin{aligned} A_{i,k}e^{p_i - \eta_{i,k} + \zeta_{i,k}} &= \text{mid}(A_{i,k}e^{p_i}, U_{i,k}e^{p_k}, L_{i,k}e^{p_k}) \\ A_{k,j}e^{-p_j - \eta_{k,j} + \zeta_{k,j}} &= \text{mid}(A_{k,j}e^{-p_j}, U_{k,j}e^{-p_k}, L_{k,j}e^{-p_k}) . \end{aligned}$$

This gives an expression of coordinate descent with no storage of η_{k+1} nor ζ_{k+1} .

Algorithm 6.6 Exact coordinate descent for Truncated scaling

Given an initial vector p^0 , calculate iteratively p^k by selecting a coordinate l following a cyclic rule and computing: $p_l^{k+1} = p_l^k$ for all $l' \neq l$ and

$$\begin{aligned} p_l^{k+1} &= \frac{1}{2} \log \left(\sum_{i \in [n]} \text{mid}(A_{i,l}e^{p_i^k}, U_{i,l}e^{p_l^k}, L_{i,l}e^{p_l^k}) \right) \\ &\quad - \frac{1}{2} \log \left(\sum_{j \in [n]} \text{mid}(A_{l,j}e^{-p_j^k}, U_{l,j}e^{-p_l^k}, L_{l,j}e^{-p_l^k}) \right) \end{aligned}$$

Proposition 6.10. *If A has a truncated scaling, then Algorithm 6.6 converges linearly to a solution of the truncated scaling problem.*

Proof. By Proposition 6.9, we can see that Algorithm 6.6 is a coordinate descent algorithm (Algorithm 6.3) applied to the minimization of the dual function θ defined in (6.9) such that for every l , the coordinate selection order is

$$\eta_{l,1}, \dots, \eta_{l,n}, \zeta_{l,1}, \dots, \zeta_{l,n}, p_l$$

As in Proposition 6.8, we then just verify the hypotheses of Proposition 6.1. \square

Remark 6.6. Due to the truncation, it is not clear how to determine the primitivity of the gradient of the fixed point operator of a HOTS-like algorithm for the truncated scaling problem.

Tomlin proposed in [Tom03] to search for a flow of webservers ρ that maximizes the entropy for the effective HOTS problem (Section 6.3) with additional bound constraints of the type

$$U_{i,j} \leq \rho_{i,j} \leq L_{i,j} . \quad (6.10)$$

These bound constraints model the fact that one may have information on the actual flow of webservers through some hyperlink, even if the flow on every hyperlink is out of reach.

We propose the following coordinate descent algorithm (Algorithm 6.7) that couples the algorithms presented in Proposition 6.8 and Algorithm 6.6. The proof of convergence is just the concatenation of the arguments of Propositions 6.8 and 6.10.

Algorithm 6.7 Coordinate descent for the HOTS problem with bounds

Start with an initial point $y^0 \in \mathbb{R}^n$, $y^0 > 0$. Given y^k , select a coordinate $i \in [n + 1]$ and compute y^{k+1} such that $y_j^{k+1} = y_j^k$, $\forall j \neq i$ and

$$y_i^{k+1} = \left(\frac{\sum_{i \in [n]} \text{mid}(A_{i,l} y_i^k, U_{i,l} y_l^k, L_{i,l} y_l^k) + e^{a^k} y_{n+1}^k}{\sum_{j \in [n]} \text{mid}(A_{l,j} (y_j^k)^{-1}, U_{l,j} (y_l^k)^{-1}, L_{l,j} (y_l^k)^{-1}) + e^{-b^k} (y_{n+1}^k)^{-1}} \right)^{\frac{1}{2}}$$

If $i < n + 1$, then set $(\mu^{k+1}, a^{k+1}, b^{k+1}) = (\mu^k, a^k, b^k)$, otherwise

$$\begin{aligned} \mu^{k+1} &= \log\left(\frac{2\alpha - 1}{\sum_{j,j' \in [n]} \text{mid}(A_{j,j'} \frac{y_j^k}{y_{j'}^k}, U_{j,j'}, L_{j,j'})}\right) , \\ a^{k+1} &= \log\left(\frac{1 - \alpha}{\sum_{j \in [n]} \frac{y_{n+1}^k}{y_j^k} e^{\mu^{k+1}}}\right) , \\ b^{k+1} &= -\log\left(\frac{1 - \alpha}{\sum_{j' \in [n]} \frac{y_{j'}^k}{y_{n+1}^k} e^{\mu^{k+1}}}\right) . \end{aligned}$$

6.5 Comparison with alternative algorithms

We give in Table 6.1 a comparison of four algorithms for the matrix balancing problem: interior-reflective Newton method (Matlab `fminunc` function), coordinate descent, DomEig and ideal HOTS. We considered a small matrix, a medium size matrix and two large matrices.

The matrix $A = \begin{bmatrix} \epsilon & 1 \\ 2 & 0 \end{bmatrix}$, with $\epsilon = 10^{-3}$, is a nearly imprimitive matrix. The CMAP matrix is the adjacency matrix of a fragment of the web graph of size 1,500. The crawl consists of the `www.cmap.polytechnique.fr` website and surrounding pages. The NZ Uni matrix comes from a crawl of New Zealand Universities websites, available at [Pro06]. It has 413,639 pages. The uk2002 matrix comes from a crawl of the .uk name domain with 18,520,486 pages, gathered by UbiCrawler [BCSV04]. For the matrix balancing problem, we added to the entries of the adjacency matrices arising from fragments of the web graph a small positive constant equal to $1/n$ (to guarantee irreducibility of the matrix). We launched our numerical experiments on a personal computer with 4 Intel Xeon CPUs at 2.98 GHz and 8 GB RAM.

	$A = \begin{bmatrix} \epsilon & 1 \\ 2 & 0 \end{bmatrix}$	CMAP 1,500 p.	NZ Uni 413,639 p.	uk2002 18,520,486 p.
$ \lambda_2(P) $ (Thm. 6.2)	0.9993	0.8739	0.9774	0.998
Matlab's fminunc	0.015 s	948 s	out of mem.	-
DomEig [JPS00]	1.87 s	34.4 s	> 600 s	-
Coordinate desc. (Alg. 6.2)	0.006 s	0.03 s	6.06 s	2,391 s
Ideal HOTS (Alg. 6.1)	2.0 s	0.02 s	7.52 s	1,868 s

Table 6.1: Execution times for 4 algorithms to solve the matrix balancing problem. General purpose algorithms like Newton methods (fminunc) are outperformed by coordinate descent and ideal HOTS for this problem. DomEig [JPS00] does not seem to be very efficient for these sparse problems. On the other hand, coordinate descent and ideal HOTS have similar performances. We remark however that coordinate descent has a better behavior for imprimitive matrices but that we have the expression of the rate of convergence of ideal HOTS, that we do not have for coordinate descent.

Convex optimization solvers, coding algorithms like quasi-Newton, are heavy machineries that can reach quadratic convergence and can handle general problems. They however need complex algorithms: one should not program them by scratch but use robust public codes. Parallel computation is not so trivial and tuning the parameters may be difficult. Second order methods also need to compute the Hessian matrix, which may be a large dense matrix.

The DomEig algorithm [JPS00] consists of 2 loops. The inner loop is the power method, it has a linear speed of convergence equal to the spectral gap of the matrix considered. Most of the computational cost lies in the power iterations, since the outer loop is simple. This algorithm is specially designed for matrix balancing and accepts no extension. It does not seem to be very efficient for sparse problem, when compared to HOTS and coordinate descent. Moreover, the precision required for the determination of the principal eigenvalue has a strong impact on the performance of the algorithm.

Schneider's coordinate descent DSS algorithm for the matrix balancing problem is a very efficient and scalable algorithm. It remains efficient for non primitive matrices.

The ideal HOTS algorithm has a linear speed of convergence equal to the spectral gap of P . On our experiments, it performs well for medium and large size problems but lacks efficiency for imprimitive matrices. For the fragments of the web graph, it converged in a bit more iterations than coordinate descent but each iteration is a bit simpler. So the execution times are about the same. Its advantage compared to coordinate descent is that its implementation only needs left and right matrix-vector products and element-wise operations (division and square root). Hence, it does not require to compute and store the transpose of the adjacency matrix, which can be useful for web scale applications.

In Table 6.2, we compare coordinate descent and the effective HOTS algorithm for the effective HOTS problem. Both algorithms scale well and have comparable computational costs. However, we remark that the rate of convergence seems to deteriorate when the size of the problem increases. It might be necessary to choose smaller values for α for large graphs in order to compensate this phenomenon. In the next section, we propose an alternative solution consisting in normalizing the adjacency matrix prior to computing the HOTS score, and hence minimizing a relative entropy instead of the entropy function.

In Table 6.3, we give the execution times for the HOTS problem with bounds (6.10).

	$A = \begin{bmatrix} \epsilon & 1 \\ 2 & 0 \end{bmatrix}$	CMAP	NZ Uni	uk2002
$ \lambda_2(\nabla F) $ (Prop. 6.5)	0.8846	0.946	0.995	0.9994
Coordinate descent (Prop. 6.8)	0.0061 s	0.0613 s	35.7 s	3,809 s
Effective HOTS (Alg. 6.5)	0.0217 s	0.0589 s	36.5 s	3,017 s
PageRank [BP98]	0.0216 s	0.0195 s	2.9 s	270 s

Table 6.2: Comparison of coordinate descent and HOTS for the effective HOTS problem with $\alpha = 0.9$. Both algorithms seem to perform equally. In particular, unlike PageRank, the convergence rate deteriorates with the size of the fragment of the web graph considered.

	$A = \begin{bmatrix} \epsilon & 1 \\ 2 & 0 \end{bmatrix}$	CMAP	NZ Uni
Inexact coor. descent (Eq. (6.8) and [Sch90])	0.019 s	4.28 s	> 2000 s
Exact coor. descent (Alg. 6.7)	0.005 s	0.23 s	100 s

Table 6.3: Execution times for two algorithms to solve the effective HOTS problem where some arcs have prescribed bounds on their flow [Tom03] (these bounds are determined at random). Exact coordinate descent (Algorithm 6.7) is faster than inexact coordinate descent. Indeed we do not need any line search for Algorithm 6.7. By comparison with Table 6.2, we can see that the computational expense is multiplied by less than 4 with the bound constraints.

6.6 A normalized HOTS algorithm for a better convergence rate

The previous study has shown two drawbacks of Tomlin's HOTS algorithm. First of all, its convergence rate seems to deteriorate when the size of the problem increases. The second problem concerns manipulations of the HOTS score: when one single page is considered, a very good strategy is to point to no page, and thus make this page a dangling node in the web graph. This comes from the relation of HOTS with the anti-Perron ranking (Remark 6.3). Indeed, the anti-Perron ranking penalizes bad quality links but also adding any outlink, even a good quality one, diminishes the score of the page where it has been added.

We now propose a modification of the HOTS algorithm that tackles those two issues. In order to stop penalizing the presence of hyperlinks on a page, we normalize the adjacency matrix, and thus state the entropy optimization problem as a relative entropy optimization problem. Then the Perron ranking reduces to PageRank and the normalized anti-Perron ranking does not penalize the number of links any more.

We also need to address the problem of dangling nodes, in which the normalization of the corresponding row of the adjacency matrix is not defined, and the reducibility of the adjacency matrix. A possibility is to keep on inspiring from the PageRank and consider the Google matrix instead of the normalized adjacency matrix. With this choice, dangling pages are considered to point to every page, which implies that they have a low rank in the normalized HOTS score, when compared to the rank given by PageRank. Instead, we suggest to set Tomlin's effective network model to solve the reducibility problem and add another fictitious node that point to every page and is pointed to by every dangling page.

We end up with the following network flow problem where the $(n + 2) \times (n + 2)$ matrix

M is defined by:

$$M_{i,j} = \begin{cases} \frac{A_{i,j}}{\sum_k A_{i,k}} & \text{if } i, j \leq n, \sum_k A_{i,k} \geq 1 \\ 0 & \text{if } i, j \leq n, \sum_k A_{i,k} = 0 \end{cases}$$

$$M = \begin{bmatrix} M_{[n],[n]} & f & 1 \\ 1^T & 0 & 1 \\ 1^T & 1 & 0 \end{bmatrix}$$

and f is the 0-1 vector such that $f_i = 1$ if and only if i is a dangling node.

$$\max_{\rho \geq 0} - \sum_{i,j \in [n+2]} \rho_{i,j} \left(\log \left(\frac{\rho_{i,j}}{M_{i,j}} \right) - 1 \right) \quad (6.11)$$

$$\sum_{j \in [n+2]} \rho_{i,j} = \sum_{j \in [n+2]} \rho_{j,i}, \quad \forall i \in [n+2] \quad (p_i)$$

$$\sum_{i,j \in [n+2]} \rho_{ij} = 1 \quad (\mu)$$

$$\sum_{j \in [n]} \rho_{n+2,j} = 1 - \alpha \quad (a)$$

$$1 - \alpha = \sum_{i \in [n]} \rho_{i,n+2} \quad (b)$$

We call this optimization problem the normalized HOTS problem. The normalized HOTS algorithm is defined in the same way as the effective HOTS algorithm but with M instead of A' . As in Lemma 6.1, we define

$$\mu' = \log \left(\frac{2\alpha - 1}{\sum_{i,j \in [n]} M_{i,j} e^{p_i - p_j}} \right),$$

$$a' = \log \left(\frac{1 - \alpha}{2\alpha - 1} \frac{\sum_{i,j \in [n]} M_{i,j} e^{p_i - p_j}}{\sum_{j \in [n]} e^{p_{n+1} - p_j}} \right),$$

$$b' = -\log \left(\frac{1 - \alpha}{2\alpha - 1} \frac{\sum_{i,j \in [n]} M_{i,j} e^{p_i - p_j}}{\sum_{i \in [n]} e^{p_i - p_{n+1}}} \right)$$

and $\lambda'(p) = (\mu', a', b')$. For $\lambda' \in \mathbb{R}^3$, we denote

$$g_i^{\lambda'}(y) = \left(\frac{\sum_{j \in [n]} M_{j,i} y_j + e^{a'} y_{n+1}}{\sum_{k \in [n]} M_{i,k} (y_k)^{-1} + e^{-b'} (y_{n+1})^{-1}} \right)^{\frac{1}{2}}.$$

Algorithm 6.8 Normalized HOTS algorithm

Given an initial point $y^0 \in \mathbb{R}^n$, the effective HOTS algorithm is defined by

$$y^{k+1} = g^{\lambda'(\log(y^k))}(y^k)$$

	CMAP	NZ Uni	uk2002
$ \lambda_2(\nabla F) $ (classical HOTS)	0.946	0.995	0.9994
$ \lambda_2(\nabla F) $ (normalized HOTS)	0.906	0.988	0.96
Execution times for Normalized HOTS	0.055 s	46.24 s	752 s

Table 6.4: Performances of the normalized HOTS algorithm presented in Section 6.6. The correlation between the deterioration of the convergence rate of the algorithm, given by the second eigenvalue of the matrix ∇F , and the size of the data set does not hold any more. Moreover, on all the tests we performed, the convergence rate remained under 0.99.

Proposition 6.11. *If there exists a primal feasible point to (6.11) with the same pattern as A' , then the normalized HOTS algorithm converges with a linear rate of convergence.*

Proof. In the proof of Theorem 6.5 we did not use the fact that the adjacency matrix A is a 0-1 matrix, only that its elements are nonnegative. Hence, the convergence proof of the effective HOTS algorithm directly applies to the normalized HOTS algorithm. \square

We shall see in Table 6.4 that the convergence properties of this algorithm seem to be better than those of the classical HOTS algorithm. Nevertheless, these experimental results are still to be validated by other theoretical and numerical studies.

CHAPTER 7

Optimization of the Perron eigenvalue and eigenvector

7.1 Introduction

Motivation The main problem of this chapter is the optimization of the ranking of a given web site. It consists in finding an optimal outlink strategy maximizing a given ranking subject to design constraints and for a given ranking algorithm.

One of the main ranking methods relies on the PageRank introduced by Brin and Page [BP98]. We studied the PageRank optimization problem in Chapter 4. In the present chapter, that follows [Fer12b], we consider the more general situation in which the ranking is determined by the Perron eigenvector of a nonnegative, but not necessarily stochastic, matrix. The Perron-Frobenius theorem (see [BP94] for instance) states that any nonnegative matrix A has a nonnegative principal eigenvalue called the Perron root and nonnegative principal eigenvectors. If, in addition, A is irreducible, then the Perron root is simple and the (unique up to a multiplicative constant) nonnegative eigenvector, called the Perron vector, has only positive entries. This property makes it a good candidate to sort web pages. The ranking algorithms considered differ in the way of constructing from the web graph a nonnegative irreducible matrix from which we determine the Perron vector. Then, the greater is the Perron vector's coordinate corresponding to a web page, the higher this web page is in the ranking. In [Kee93], such a ranking is proposed for football teams. The paper [Saa87] uses the Perron vector to rank teachers from pairwise comparisons. See also [Vig09b] for a survey on the

subject. When it comes to web page ranking, the PageRank is the Perron eigenvector of the transition matrix described above but the HITS algorithm also rank pages according to a Perron vector. More examples of Perron vector rankings or nonlinear Perron vector rankings for the web, like the Sinkhorn ranking, are given in Chapter 2.

The HITS algorithm [Kle99] is not purely a link-based algorithm. It is composed of two steps and the output depends on the query of the user. Given a query, we first select a seed of pages that are relevant to the query according to their text content. This seed is then extended with pages linking to them, pages to which they link and all the hyperlinks between the pages selected. We thus obtain a subgraph of the web graph focused on the query. Then, the second step assigns each page two scores: a hub score v and an authority score u such that good hubs should point to good authorities and good authorities should be pointed to by good hubs. Introducing the adjacency matrix A of the focused graph, this can be written as $v = \rho Au$ and $u = \rho A^T v$ with $\rho \in \mathbb{R}_+$, which means that the vector of hub scores is the Perron eigenvector of the matrix $A^T A$ and that the vector of authority scores is the Perron eigenvector of AA^T . The construction of HITS's focused subgraph is a combination of text content relevancy with the query and of hyperlink considerations. Maximizing the probability of appearance of a web page on this subgraph is thus a composite problem out of the range of this work. We shall however study the optimization of HITS authority, for a given focused subgraph.

We shall also study the optimization of Tomlin's HOTS scores [Tom03]. In this case, the ranking is the vector of dual variables of an optimal flow problem. The flow represents an optimal distribution of web surfers on the web graph in the sense of entropy minimization. The dual variable, one by page, is interpreted as the "temperature" of the page, the hotter a page the better. Tomlin showed that this vector is solution of a nonlinear fixed point equation: it may be seen as a nonlinear eigenvector but the fixed point operator is not necessarily monotone. However, we show that most of the arguments available in the case of Perron vector optimization can be adapted to HOTS optimization.

Contribution In this chapter, we study the problem of optimizing the Perron eigenvector of a controlled matrix and apply it to HITS and HOTS optimization. Our first main result is the development of a scalable algorithm for the local optimization of a scalar function of the Perron eigenvector over a set of nonnegative irreducible matrices. Indeed, the global Perron vector optimization over a convex set of nonnegative matrices is NP-hard, so we focus on the searching of local optima. We give in Theorem 7.1 a power-type algorithm for the computation of the matrix of the partial derivatives of the objective, based on the fact that it is a rank 1 matrix. This theorem shows that computing the partial derivatives of the objective has the same order of complexity as computing the Perron vector by the power method, which is the usual method when dealing with the large and sparse matrices built from the web graph. Then we give an optimization algorithm that couples power and gradient iterations (Algorithms 7.2 and 7.3). Each step of the optimization algorithm involves a suitable number of power iterations and a descent step. By considering this algorithm to be an approximate projected gradient algorithm [Pol97, PP02], we prove that the algorithm converges to a stationary point (Theorem 7.2). Compared with the case when the number of power iterations is not adapted dynamically, we got a speedup between 3 and 20 in our numerical experiments (Section 7.8) together with a more precise convergence result.

Our second main result is the application of Perron vector optimization to the optimiza-

tion of scalar functions of HITS authority or other scores like Tomlin's HOTS. We derive optimization algorithms and, thanks to the low rank of the matrix of partial derivatives, we show that the optimal linkage strategies of both problems satisfy a threshold property (Propositions 7.11 and 7.14). This property was already proved for PageRank optimization in [dKNvD08, FABG13] (see Propositions 2.3 and 4.11 and Theorems 4.4 and 4.5). As in [IT09, CJB10, FABG13] we partition the set of potential links (i, j) into three subsets, consisting respectively of the set of *obligatory links*, the set of *prohibited links* and the set of *facultative links*. When choosing a subset of the facultative links, we get a graph from which we get any of the three ranking vectors. We are then looking for the subset of facultative links that maximizes a given utility function. We also study the associated relaxed problems, where we accept weighted adjacency matrices. This assumes that the webmaster can influence the importance of the hyperlinks of the pages she controls, for instance by choosing the size of the font, the color or the position of the link within the page. In fact, we shall solve the relaxed problems and then give conditions or heuristics to get an admissible strategy for the discrete problems.

Related works As explained in the first part of the introduction, this work extends the study of PageRank optimization developed in [AL06, MV06, dKNvD08, IT09, CJB10, FABG13] to HITS authority [Kle99] and HOTS [Tom03] optimization.

We based our study of Perron eigenvector optimization on two other domains: eigenvalue optimization and eigenvector sensitivity. There is a vast literature on eigenvalue and eigenvector sensitivity with many domains of application (see the survey [HA89] for instance). These works cope with perturbations of a given system. They consider general square matrices and any eigenvalue or eigenvector. They give the directional derivatives of the eigenvalue and eigenvector of a matrix with respect to a given perturbation of this matrix [Nel76, MS88]. Perron eigenvalue and eigenvector sensitivity was developed in [DN84, DN85].

This led to the development of eigenvalue optimization. In [CDW75, Ove91, SF95] the authors show that the minimization of a convex function of the eigenvalues of symmetric matrices subject to linear constraints is a convex problem and can be solved with semi-definite programming. Eigenvalue optimization of nonsymmetric matrices is a more difficult problem. In general, the eigenvalue is a nonconvex nonlipschitz function of the entries of the matrix. The last section of [LO96] proposes a method to reduce the nonsymmetric case to the symmetric case by adding (many) additional variables. Another approach is developed in [OW88]: the authors derive descent directions and optimality conditions from the perturbation theory and uses so-called dual matrices.

The optimization of the Perron eigenvalue, which is equal to the spectral radius, over a set of nonnegative matrices, has recently motivated several papers. In [HNT99], Han et al. study the maximization of the spectral radius of a nonnegative matrix subject to fixed Frobenius norm perturbations and they show that the optimal perturbation has a rank 1. Axtell et al. [AHH⁺09] study the maximization or minimization of the Perron eigenvalue over a set of matrices of the form $M = SA$ where A is a fixed nonnegative matrix and S is a doubly stochastic matrix. They show that in both cases the optimum is attained at a permutation matrix. In [BS07], Boche and Schubert give an algorithm for the minimization of the Perron eigenvalue over a set of matrices with independently controlled columns. Nesterov and Protasov [NP11] generalized this result to the minimization and maximization problems. They based their approach on a convex relaxation of the Perron eigenvalue optimization problem

that they showed to be exact for product families of nonnegative operators (a generalization of matrices with independently controlled columns). Their theorem also applies to the joint spectral radius, a quantity linked to the maximal spectral radius and that is also difficult to compute in general [TB97, BN05].

In the context of population dynamics, the problem of the maximization of the growth rate of a population can be modeled by the maximization of the Perron value of a given family of matrices. This technique is used in [Log08] to identify the parameters of a population dynamic model, in [BCF⁺11a] for chemotherapy optimization purposes. An approach based on branching Markov decision processes is presented in [RW82]. Perron value optimization also appears in other contexts like in the minimization of the interferences in a network [BS07].

Apart from the stochastic case which can be solved by Markov decision techniques, like for PageRank, the search for a matrix with optimal eigenvectors does not seem to have been much considered in the literature. Indeed, the problem is not well defined since when an eigenvalue is not simple, the associated eigenspace may have dimension greater than 1.

Organization The chapter is organized as follows. In Section 7.2, we introduce Perron eigenvector and eigenvalue optimization problems and show that these problems are NP-hard problems on convex sets of matrices. We also point out some problems solvable in polynomial time. In Section 7.3, we give in Theorem 7.1 a method for the efficient computation of the derivative of the objective function. Then in Section 7.4, we give the coupled power and gradient iterations and its convergence proof. In Section 7.5, we show how HITS authority optimization problems and HOTS optimization problems reduce to Perron vector optimization. In Section 7.6, we extend these results to the optimization of a scalar function of the eigenvector of a monotone homogeneous map and in Section 7.7, we apply it to the optimization of HOTS. Finally, we report numerical results on a fragment of the web graph in Section 7.8.

7.2 Perron vector and Perron value optimization problems

Let $M \in \mathbb{R}^{n \times n}$ be a (elementwise) nonnegative matrix. We say that M is irreducible if it is not similar to a block upper triangular matrix with two blocks via a permutation. Equivalently, define the directed graph with n nodes and an edge between node i and j if and only if $M_{i,j} > 0$: M is irreducible if and only if this graph is strongly connected.

We denote by $\rho(M)$ the principal eigenvalue of the irreducible nonnegative matrix M , called the Perron root. By Perron-Frobenius theorem (see [BP94] for instance), we know that $\rho(M) > 0$ and that this eigenvalue is simple. Given a normalization N , we denote by $u(M)$ the corresponding normalized eigenvector, called the Perron vector. The normalization is necessary since the Perron vector is only defined up to positive multiplicative constant. The normalization function N should be homogeneous and we require $u(M)$ to verify $N(u(M)) = 1$. The Perron-Frobenius theorem asserts that $u(M) > 0$ elementwise. The Perron eigenvalue optimization problem on the set \mathcal{M} can be written as:

$$\min_{M \in \mathcal{M}} f(\rho(M)) \quad (7.1)$$

The Perron vector optimization problem can be written as:

$$\min_{M \in \mathcal{M}} f(u(M)) \quad (7.2)$$

We assume that f is a real valued continuously differentiable function; \mathcal{M} is a set of irreducible nonnegative matrices such that $\mathcal{M} = h(\mathcal{C})$ with h continuously differentiable and \mathcal{C} a closed convex set. These hypotheses allow us to use algorithms such as projected gradient for the searching of stationary points.

We next observe that the minimization of the Perron root and the optimization of a scalar function of the Perron vector are NP-hard problems and that only exceptional special cases appear to be polynomial time solvable by current methods. Consequently, we shall focus on the local resolution of these problems, with an emphasis on large sparse problems like the ones encountered for web applications. We consider the two following problems:

PERRONROOT_MIN: given a rational linear function $A : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ and a vector b in \mathbb{Q}^m , find a matrix M that minimizes $\rho(M)$ on the polyhedral set $\{M \in \mathbb{R}^{n \times n} \mid A(M) \leq b, M \geq 0\}$.

PERRONVECTOR_OPT: given a rational linear function $A : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$, a vector b in \mathbb{Q}^m , a rational function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and a rational normalization function N , find a matrix M that minimizes $f(u(M))$ on the polyhedral set $\{M \in \mathbb{R}^{n \times n} \mid A(M) \leq b, M \geq 0\}$, where $u(M)$ verifies $N(u(M)) = 1$.

In general, determining whether all matrices in an interval family are stable is a NP-hard problem [BT00]. The corresponding problem for nonnegative matrices is to determine whether the maximal Perron root is smaller than a given number. Indeed, as the Perron root is a monotone function of the entries of the matrix (see Proposition 7.6 below), this problem is trivial on interval matrix families. However, we shall prove NP-hardness of **PERRONROOT_MIN** and **PERRONVECTOR_OPT** by reduction of linear multiplicative programming problem to each of these problems. The linear multiplicative problem is:

LMP: given a $n \times m$ rational matrix A and a vector b in \mathbb{Q}^m , find a vector $x \in \mathbb{R}^n$ that minimizes $x_1 x_2$ on the polyhedral set $\{x \in \mathbb{R}^n \mid Ax \leq b, x_1, x_2 \geq 0\}$.

A theorem of Matsui [Mat96] states that **LMP** is NP-hard. We shall need a slightly stronger result about a weak version of **LMP**:

Weak-LMP: given $\epsilon > 0$, a $n \times m$ rational matrix A and a vector b in \mathbb{Q}^m , find a vector $x \in \mathbb{Q}$ such that $x_1 x_2 \leq y_1 y_2 + \epsilon$ for all $y \in Q$, where $Q = \{x \in \mathbb{R}^n \mid Ax \leq b, x_1, x_2 \geq 0\}$.

Lemma 7.1. *Weak-LMP is a NP-hard problem.*

Proof. A small modification of the proof of Matsui [Mat96] gives the result. If we replace $g(x_0, y_0) \leq 0$ by $g(x_0, y_0) \leq -2$ in Corollary 2.3 we remark that the rest of the proof still holds since $n^4 p^{4n} + p^2 - 4p^{4n+1} \leq -2$ for all $n \geq 1$ and $p = n^{n^4}$. Then, with the notations of [Mat96], we have proved that the optimal value of $P1(M)$ is less than or equal to $4p^{8n} - 2$ if and only if $Mx = 1$ has a 0 – 1 valued solution and it is greater than or equal to $4p^{8n}$ if and only if $Mx = 1$ does not have a 0 – 1 valued solution. We just need to choose $\epsilon < 2$ in problem $P1(M)$ to finish the proof of the lemma. \square

Proposition 7.1. *PERRONROOT_MIN and PERRONVECTOR_OPT are NP-hard problems.*

Proof. We define the matrix M by the matrix with lower diagonal equal to x and 1 on the

top left corner:

$$M = \begin{pmatrix} 0 & 0 & \dots & 0 & 1 \\ x_1 & 0 & 0 & & 0 \\ 0 & x_2 & 0 & & \vdots \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & x_n & 0 \end{pmatrix}. \quad (7.3)$$

We set the admissible set X for the vector x as $X = \{x \in \mathbb{R}^m \mid Ax \geq b, x \geq 0\}$ with a rational $p \times m$ matrix A and a rational vector b of length p .

For the eigenvector problem, we set the normalization $u_1 = 1$ and we take $f(u) = u_n$. We have $u_n = \rho$, so the complexity is the same for eigenvalue and eigenvector optimization in this context.

Now minimizing $\rho(M)$ is equivalent to minimizing $x_1 x_2 \dots x_n$ because the n -th root is an nondecreasing function on \mathbb{R}_+ . We thus just need to reduce in polynomial time the ϵ -solvability of weak-LMP to the minimization of $x_1 x_2 \dots x_n$ on X .

As $x \mapsto \log(x_1 x_2 \dots x_n)$ is a concave function, either the problem is unbounded or there exists an optimal solution that is an extreme point of the polyhedral admissible set (Theorem 3.4.7 in [BS06]). Hence, using Lemma 6.2.5 in [GLS88] we see that we can define a bounded problem equivalent to the unbounded problem and with coefficients that have a polynomial encoding length: in the following, we assume that the admissible set is bounded.

Given a rational number $\epsilon > 0$, a rational $p \times m$ matrix A' and a rational vector b' of length p , let $X' := \{x \in \mathbb{R}^m \mid A'x \geq b', x_1 \geq 0, x_2 \geq 0\}$ be a bounded polyhedron. Denoting $v_2 := \min_{x \in X'} x_1 x_2$, we are looking for $x' \in X'$ such that $|y_1 y_2 - v_2| \leq \epsilon$.

Compute $\underline{C} := \min_{i \in [n]} \min_{x \in X'} x_i$ and $\bar{C} := \max_{i \in [n]} \max_{x \in X'} x_i$ (linear programs). We first set $m_0 = \bar{C} - \underline{C} + 1$ so that $m_0 > 0$ and $m_0 > -\underline{C}$ and $t^{m_0} \in \mathbb{R}^n$ defined by $t_i^{m_0} = 0$ if $i \in \{1, 2\}$ and $t_i^{m_0} = m$ if $i \geq 3$. Let

$$X_0 := \{x \in \mathbb{R}^m \mid A'x \geq b' - A't^{m_0}, x \geq 0\}.$$

We have $v_2 = \min_{x \in X'} x_1 x_2 = \min_{x \in X_0} x_1 x_2$. Let $v_0 := \min_{x \in X} x_1 x_2 \dots x_n$. For all $x \in X_0$, we have

$$x_1 x_2 (m_0 + \underline{C})^{n-2} \leq x_1 x_2 \dots x_n \leq x_1 x_2 (m_0 + \bar{C})^{n-2},$$

so that $v_2 (m_0 + \underline{C})^{n-2} \leq v_0 \leq v_2 (m_0 + \bar{C})^{n-2}$.

We now set

$$m := \max\left\{-\underline{C} + 2^{n-3} \frac{\bar{C}}{\epsilon} \frac{v_0}{(m_0 + \underline{C})^{n-2}}, \bar{C} - 2\underline{C}, 1\right\}$$

and we define t^m and X in the same way as t^{m_0} and X_0 . Remark that m an encoding length polynomial in the length of the entries. Let $v_n := \min_{x \in X} x_1 x_2 \dots x_n$. For all $x \in X$, we have $v_2 (m + \underline{C})^{n-2} \leq v_n \leq v_2 (m + \bar{C})^{n-2}$ and $x' = x - t^m$ is a point of X' with $x'_1 x'_2 = x_1 x_2$.

As $v_2 (m_0 + \underline{C})^{n-2} \leq v_0$, $m \geq -\underline{C} + 2^{n-3} \frac{\bar{C}}{\epsilon} v_2$. As $m \geq \bar{C} - 2\underline{C}$, $\frac{m + \bar{C}}{m + \underline{C}} \geq \frac{1}{2}$, so that $\frac{(m + \underline{C})^{n-2}}{(m + \bar{C})^{n-3}} \geq \frac{1}{2^{n-3}} (m + \underline{C})$ and $\frac{(m + \underline{C})^{n-2}}{(m + \bar{C})^{n-3}} \geq \frac{\bar{C}}{\epsilon} v_2$.

Denote $M := (m + \underline{C})^{n-2}$ and $\Delta := (m + \bar{C})^{n-2} - (m + \underline{C})^{n-2}$: $M v_2 \leq v_n \leq (M + \Delta) v_2$. We have $\Delta := \sum_{k=0}^{n-3} \binom{n-2}{k} m^k \bar{C}^{n-2-k} - \sum_{k=0}^{n-3} \binom{n-2}{k} m^k \underline{C}^{n-2-k} \leq \bar{C} (m + \bar{C})^{n-3}$. Hence, $M \geq \frac{\Delta}{\epsilon} v_2$. As $\Delta \geq 0$, $M + \Delta \geq \frac{\Delta}{\epsilon} v_2$. We obtain

$$\epsilon \geq \frac{\Delta M v_2}{M(M + \Delta)} = \left(\frac{1}{M} - \frac{1}{M + \Delta}\right) M v_2 \geq v_2 - \frac{v_n}{M + \Delta}.$$

Finally

$$\frac{v_n}{(m_0 + \bar{C})^{n-2}} \leq v_2 \leq \frac{v_n}{(m_0 + \bar{C})^{n-2}} + \epsilon .$$

which proves that weak-LMP reduces to the minimization of $x_1 x_2 \dots x_n$ on X . \square

An alternative reduction is given in [Vla12]. There, the author proves that minimizing the Perron eigenvalue over a convex set defined by its extreme points is NP-hard. The reduction is complementary to ours since in Proposition 7.1, we define the convex set by the equations of its facets. Note that both reductions do not extend to Perron vector maximization.

The general Perron eigenvalue optimization problem is NP-hard but we however point out some cases for which it is tractable. The following proposition is well known:

Proposition 7.2 ([Kin61]). *The eigenvalue $\rho(M)$ is a log-convex function of the log of the entries of the nonnegative matrix M .*

This means that $\log \circ \rho \circ \exp$ is a convex function, where \exp is the componentwise exponential, namely if $0 \leq \alpha \leq 1$ and A and B are two nonnegative $n \times n$ matrices then for $C_{i,j} = A_{i,j}^\alpha B_{i,j}^{1-\alpha}$, $\rho(C) \leq \rho(A)^\alpha \rho(B)^{1-\alpha}$.

Corollary 7.1. *The optimization problem*

$$\min_{M \in \exp(C)} \rho(M)$$

with C convex is equivalent to the convex problem

$$\min_{L \in C} \log \circ \rho \circ \exp(L) .$$

The difference between this proposition and the previous one is that here $\mathcal{M} = h(C) = \exp(C)$ whereas previously we had h affine. This makes a big difference since an ϵ -solution of a convex program can be found in polynomial time [BTN01].

Remark 7.1. The largest singular value (which is a norm) is a convex function of the entries of the matrix. For a symmetric matrix, the singular values are the absolute values of the eigenvalues. Thus minimizing the largest eigenvalue on a convex set of nonnegative symmetric matrices is a convex problem.

In order to solve the signal to interference ratio balancing problem, Boche and Schuber [BS07] give an algorithm for the global minimization of the Perron root when the rows of the controlled matrix are independently controlled, ie when the admissible set is of the form $\mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ and $z_k \in \mathcal{Z}_k$ is the k^{th} row of the matrix.

Proposition 7.3 ([BS07]). *Let $\mathcal{Z} = \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ and Γ be a fixed positive diagonal matrix. If the k^{th} row of the matrix $V(z)$ only depends on $z_k \in \mathcal{Z}_k$, $V(z)$ is irreducible for all $z \in \mathcal{Z}$ and if $V(z)$ is continuous on \mathcal{Z} (\mathcal{Z} can be discrete), then there exists a monotone algorithm that minimizes $\rho(\Gamma V(z))$ over \mathcal{Z} , in the sense that $\rho(\Gamma V(z_{n+1})) \leq \rho(\Gamma V(z_n))$ for all $n \geq 0$ and $\lim_n \rho(\Gamma V(z_n)) = \min_{z \in \mathcal{Z}} \rho(\Gamma V(z))$.*

Nesterov and Protasov generalized this result using the following relaxations of the minimal and maximal eigenvalue.

Proposition 7.4 ([NP11]). *Let \mathcal{M} be a set of nonnegative matrices with a common invariant cone K (one often chooses $K = \mathbb{R}^n$), let $e \in \text{int}K$ and $e_* \in \text{int}K^*$ be two vectors. Denote $\Delta = \{x \in K \mid \langle e_*, x \rangle = 1\}$, $\Delta^* = \{s \in K^* \mid \langle s, e \rangle = 1\}$ and*

$$\begin{aligned}\lambda^*(\mathcal{M}) &= \inf_{\lambda \in \mathbb{R}, x \in \text{int}K} \{\lambda \mid \lambda x - Ax \in K, \forall A \in \mathcal{M}\} \\ \lambda_*(\mathcal{M}) &= \sup_{\lambda \in \mathbb{R}, x \in K} \{\lambda \mid Ax - \lambda x \in K, \forall A \in \mathcal{M}\} .\end{aligned}$$

We have the relation

$$\lambda_*(\mathcal{M}) \leq \min_{A \in \mathcal{M}} \rho(A) \leq \max_{A \in \mathcal{M}} \rho(A) \leq \lambda^*(\mathcal{M})$$

Let

$$\begin{aligned}\psi^*(x, \lambda) &= \max_{s \in \Delta^*} \max_{A \in \mathcal{M}} \langle s, Ax - \lambda x \rangle , & \xi^*(\lambda) &= \inf_{v \in \Delta} \psi^*(v, \lambda) \\ \psi_*(x, \lambda) &= \max_{s \in \Delta^*} \max_{A \in \mathcal{M}} \langle s, \lambda x - Ax \rangle , & \xi_*(\lambda) &= \inf_{v \in \Delta} \psi_*(v, \lambda)\end{aligned}$$

The functions ψ^* and ψ_* are convex (computable if $\text{co}(\mathcal{M})$ is an effective convex set for instance) and $\lambda_*(\mathcal{M})$ and $\lambda^*(\mathcal{M})$ are the zeros of the monotone functions ξ_* and ξ^* .

Finally, suppose that \mathcal{M} is a product family, that is $\mathcal{M} = \{A = BF \mid Fe_j \in \mathcal{F}_j, j \in [n]\}$ for B a fixed matrix and $\mathcal{F}_j \subset K^*$, $j \in [n]$. Then the relaxation is exact: $\lambda_*(\mathcal{M}) = \min_{A \in \mathcal{M}} \rho(A)$ and $\max_{A \in \mathcal{M}} \rho(A) = \lambda^*(\mathcal{M})$

This proposition also applies to the joint spectral radius $\sigma^*(\mathcal{M})$, a quantity linked to the maximal spectral radius and that is also difficult to compute in general [BN05], because $\max_{A \in \mathcal{M}} \rho(A) \leq \sigma^*(\mathcal{M}) \leq \lambda^*(\mathcal{M})$.

The next proposition shows that if the rows are independently controlled, one can also solve the maximum and minimum Perron eigenvalue problems thanks to nonlinear Perron-Frobenius theory (Section 2.4).

Proposition 7.5. *Let $\mathcal{Z} = \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ be a set of nonnegative matrices. Let T be the operator defined for all $x \geq 0$ and for all $i \in [n]$ by*

$$T_i(x) = \max_{z \in \mathcal{Z}_i} z \cdot x .$$

The map T is monotone, homogeneous and semi-differentiable and if there exists λ and $x > 0$ such that $T(x) = \lambda x$, then $\lambda = \max_{Z \in \mathcal{Z}} \rho(Z)$.

Moreover, if \mathcal{Z} is a compact set of primitive matrices, then the eigenvector x of T is unique and the power algorithm converges to x with a linear rate of convergence.

Similar results hold for T' defined by $T'_i(x) = \min_{z \in \mathcal{Z}_i} z \cdot x$.

Proof. As for all $i \in [n]$ and for all $z \in \mathcal{Z}_i$, the map $(x \mapsto z \cdot x)$ is monotone, homogeneous and semi-differentiable, the same holds for T_i because these properties are preserved by the maximum (Theorem 3.8 in [AGN12]) and hence for T .

Now suppose that there exist λ and $x > 0$ such that $T(x) = \lambda x$. Then for all i , there exists $y \in \mathcal{Z}_i$ such that $y \cdot x = \lambda x_i$ and for all $z \in \mathcal{Z}_i$, $y \cdot x \geq z \cdot x$. We can write it as there exists $Y \in \mathcal{Z}$ such that for all $Z \in \mathcal{Z}$, $Yx = \lambda x$ and $Zx \leq \lambda x$. We deduce using [Col42] that for all $Z \in \mathcal{Z}$, $\rho(Z) \leq \lambda = \rho(Y)$.

If \mathcal{Z} is a compact set of primitive matrices, then there exists $\alpha < 1$ such that every matrix in \mathcal{Z} has a spectral gap smaller than α . By Proposition 8.1 in [AGN11] and Theorem 6.8 in [AGN12] we deduce that T has a unique fixed point and that the power algorithm converges linearly at rate at least α . \square

When the power algorithm does not apply, one may write T in additive form as for all $i \in [n]$ and for all $X \in \mathbb{R}^n$,

$$t_i(X) = \max_{z \in \mathcal{Z}_i} \log \left(\sum_{j \in [n]} z_j e^{X_j} \right) = \max_{z \in \mathcal{Z}_i} \max_{p \in \Sigma} - \sum_j p_j \log \left(\frac{p_j}{z_j} \right) + p \cdot X ,$$

with Σ being the usual simplex (the second equality says that the relative entropy is the Fenchel transform of the log-sum-exp function). This additive formulation shows that t is the daily operator of a 0-sum game with 1 player for maximization and 2 players for minimization, and compact action spaces. Hence, it may be possible to use policy iteration on it [AD12].

We also note that other bounds on the spectral radius have been developed in the context of 0-1 matrices (adjacency matrices of graphs), like in [Fri88].

The bounds on the optimal Perron eigenvalue optimization problems are complementary with the local optimization approach developed in this chapter, which would yield an upper bound on the optimal value of this problem. However, for the Perron value optimization problems that we will study in Chapter 8, the controls do not naturally decompose line by line. Hence, the bound that we find is not tight. The following example shows that the bound can be arbitrarily bad.

Take \mathcal{M} be the set of matrices of the form (7.3) with $x \in \{y \in \mathbb{R}^n \mid y \geq 0, \sum_{i \in [n]} y_i = 1\}$. Then $\lambda^*(\mathcal{M}) = 1$ and $\max_{A \in \mathcal{M}} = n^{-\frac{n-1}{n}} \sim \frac{1}{n}$. We remark that this example is tight with respect to Theorem 1 in [BN10].

Remark 7.2. As developed in Chapter 4, general PageRank optimization problems can be formulated as follows. Let M be the transition matrix of PageRank and ρ the associated occupation measure. When M is irreducible, they are linked by $M_{ij} = \frac{\rho_{ij}}{\sum_k \rho_{ik}} = h_{ij}(\rho)$, which yields our function h . We also have $u(h(\rho))_i = \sum_k \rho_{ik}$ and $\rho_{ij} = u(M)_i M_{ij}$.

If the set \mathcal{C} , which defines the design constraints of the webmaster, is a convex set of occupation measures, if h is as defined above and if f is a convex function, then

$$\min_{\rho \in \mathcal{C}} f(u(h(\rho)))$$

is a convex problem. Thus ϵ -solutions of PageRank optimization problems can be found in polynomial-time.

To our knowledge, Perron vector optimization problems without a stochasticity assumption have scarcely been considered in the literature.

7.3 A power-type algorithm for the evaluation of the derivative of a function of the principal eigenvector

We now turn to the main topic of this chapter. We give in Theorem 7.1 a power-type algorithm for the evaluation of the partial derivatives of the principal eigenvector of a matrix with a simple principal eigenvalue.

We consider a matrix M with a simple eigenvalue λ and associated left and right eigenvectors u and v . We shall normalize v by the assumption $\sum_{i \in [n]} v_i u_i = 1$. The derivatives of the eigenvalue of a matrix are well known and easy to compute:

Proposition 7.6 ([Kat66] Section II.2.2). *Denoting v and u the left and right eigenvectors of a matrix M associated to a simple eigenvalue λ , normalized such that $\sum_{i \in [n]} v_i u_i = 1$, the derivative of λ can be written as:*

$$\frac{\partial \lambda}{\partial M_{ij}} = v_i u_j .$$

In this section, we give a scalable algorithm to compute the partial derivatives of the function $f \circ u$, that to an irreducible nonnegative matrix M associates the utility of its Perron vector. In other words we compute $g_{ij} = \sum_k \frac{\partial f}{\partial u_k} \frac{\partial u_k}{\partial M_{ij}}$. This algorithm is a sparse iterative scheme and it is the core of the optimization algorithms that we will then use for the large problems encountered in the optimization of web ranking.

We first recall some results on the derivatives of eigenprojectors (see [Kat66] for more background). Throughout the end of the chapter, we shall consider column vectors and row vectors will be written as the transpose of a column vector. Let P be the eigenprojector of M for the eigenvalue λ . One can easily check that as λ is a simple eigenvalue, the spectral projector is $P = uv^T$ as soon as $v^T u = 1$. We have the relation

$$\frac{\partial P}{\partial M_{ij}} = -SE_{ij}P - PE_{ij}S$$

where E_{ij} is the $n \times n$ matrix with all entries zero except the ij^{th} which is equal to one and $S = (M - \lambda I)^\#$ is the Drazin pseudo-inverse of $M - \lambda I$. This matrix S also satisfies the equalities

$$S(M - \lambda I) = (M - \lambda I)S = I - P \quad \text{and} \quad SP = PS = 0 . \quad (7.4)$$

When it comes to eigenvectors, we have to set a normalization for each of them. Let N be the normalization function for the right eigenvector. We assume that it is differentiable at u and that $N(\alpha u) = \alpha N(u)$ for all nonnegative scalars α , which implies $\frac{\partial N}{\partial u}(u) \cdot u = N(u)$. We normalize v by the natural normalization $\sum_i u_i v_i = 1$.

Proposition 7.7 ([MS88]). *Let M be a matrix with a simple eigenvalue λ and associated eigenvector u normalized by $N(u) = 1$. We denote $S = (M - \lambda I)^\#$. Then the partial derivatives of the eigenvector are given by:*

$$\frac{\partial u}{\partial M_{ij}}(M) = -Se_i u_j + (\nabla N(u)^T Se_i u_j)u$$

where e_i is the vector with i^{th} entry equal to 1.

To simplify notations, we denote ∇f for $\nabla f(u)$ and ∇N for $\nabla N(u)$.

Corollary 7.2. *Let M be a matrix with a simple eigenvalue λ and associated eigenvector u normalized by $N(u) = 1$. The partial derivatives of the function $M \mapsto (f \circ u)(M)$ at M are $g_{ij} = w_i u_j$, where the auxiliary vector w is given by:*

$$w^T = (-\nabla f^T + (\nabla f \cdot u)\nabla N^T)S = (-\nabla f^T + (\nabla f \cdot u)\nabla N^T)(M - \lambda I)^\#$$

Proof. By Proposition 7.7, we deduce that

$$g_{ij} = \sum_k \frac{\partial f}{\partial u_k}(u(M)) \frac{\partial u_k}{\partial M_{ij}}(M) = - \sum_k \frac{\partial f}{\partial u_k} S_{ki} u_j + \sum_k \frac{\partial f}{\partial u_k} u_k \sum_l \frac{\partial N}{\partial u_l} S_{li} u_j$$

which is the developed form the result. \square

This simple corollary already improves the computation speed. Using Proposition 7.7 directly, one needs to compute $\frac{\partial P}{\partial M_{ij}}$ in every direction, which means the computation of a Drazin inverse and $2n^2$ matrix-vector products involving M for the computation of $\left(\frac{\partial u}{\partial M_{ij}}\right)_{i,j \in [n]}$. With Corollary 7.2 we only need one matrix-vector product for the same result. The last difficulty is the computation of the Drazin inverse S . In fact, we do not need to compute the whole matrix but only to compute $S^T x$ for a given x . The next two propositions show how one can do it.

Proposition 7.8. *Let M be a matrix with a simple eigenvalue λ and associated eigenvector u normalized by $N(u) = 1$. The auxiliary vector w of Corollary 7.2 is solution of the following invertible system*

$$[w^T, w_{n+1}] \begin{bmatrix} M - \lambda I & -u \\ \nabla N^T & 0 \end{bmatrix} = [-\nabla f^T, 0] . \quad (7.5)$$

where $w_{n+1} \in \mathbb{R}$.

Proof. A nullspace argument similar to this of [Nel76] shows that $\begin{bmatrix} M - \lambda I & -u \\ \nabla N^T & 0 \end{bmatrix}$ is invertible as soon as λ is simple and $\nabla N^T u = 1$. Then the solution w of the system (7.5) verifies the equations $w^T(M - \lambda I) + w_{n+1} \nabla N^T = -\nabla f^T$ and $w^T u = 0$. Multiplying the first equality by u yields $w_{n+1} \nabla N^T u = -\nabla f^T u$ and multiplying it by S yields $w^T(I - uv^T) = -w_{n+1} \nabla N^T S - \nabla f^T S$. Putting all together, we get $w^T = (-\nabla f^T + (\nabla f^T u) \nabla N^T) S$. \square

The next proposition provides an iterative scheme to compute the evaluation of the auxiliary vector w when we consider the principal eigenvalue.

Definition 7.1. We say that a sequence $(x_k)_{k \geq 0}$ converges to a point x with a linear convergence rate α if $\limsup_{k \rightarrow \infty} \|x_k - x\|^{1/k} \leq \alpha$.

Proposition 7.9. *Let M be a matrix with only one eigenvalue of maximal modulus denoted by $\rho = |\lambda_1| > |\lambda_2|$. With the same notations as in Corollary 7.2, we denote $\tilde{M} = \frac{1}{\rho} M$ and $z^T = \frac{1}{\rho}(-\nabla f^T + (\nabla f \cdot u) \nabla N^T)$, and we fix a real row vector w_0 . Then the fixed point scheme defined by*

$$\forall k \in \mathbb{N}, \quad w_{k+1}^T = (-z^T + w_k^T \tilde{M})(I - P)$$

with $P = uv^T$, converges to $w^T = (-\nabla f^T + (\nabla f \cdot u) \nabla N^T)(M - \rho I)^\#$ with a linear rate of convergence $\frac{|\lambda_2|}{\rho}$.

Proof. We have $w_k^T = \sum_{l=0}^{k-1} -z^T (\tilde{M}(I - P))^l + w_0^T (\tilde{M}(I - P))^k$. By assumption, all the eigenvalues of \tilde{M} different from 1 have a modulus smaller than 1. Thus, using the fact that P is the eigenprojector associated to 1, we get $\rho(\tilde{M}(I - P)) = \frac{|\lambda_2|}{\rho} < 1$. By [Ost55], $(\|\tilde{M}(I - P)\|^k)^{1/k} \rightarrow \rho(\tilde{M}(I - P))$, so the algorithm converges to a limit w and for all $\epsilon > 0$,

$\|w_k - w\| = O((\frac{|\lambda_2| + \epsilon}{\rho})^k)$. This implies a linear convergence rate equal to $\frac{|\lambda_2|}{\rho}$. The limit w satisfies $w^T = (-z^T + w^T \tilde{M})(I - P)$, so $w^T P = 0$ and as $\tilde{M}P = P$, $w^T \tilde{M} - w^T = z^T(I - P)$. We thus get the equality $w^T(\tilde{M} - I) = z^T$. Multiplying both sides by $(\tilde{M} - I)^\#$, we get:

$$w^T(\tilde{M} - I)(\tilde{M} - I)^\# = w^T - w^T P = w^T = z^T(\tilde{M} - I)^\# .$$

The last equalities and the relation $(\beta^{-1}M)^\# = \beta M^\#$ show by Proposition 7.2 that $g = wu^T$ is the matrix of partial derivatives of the Perron vector multiplied by ∇f . \square

This iterative scheme uses only matrix-vector products and thus may be very efficient for a sparse matrix. In fact, it has the same linear convergence rate as the power method for the computation of the Perron eigenvalue and eigenvector. This means that the computation of the derivative of the eigenvector has a computational cost of the same order as the computation by the power method of the eigenvector itself. We next show that the eigenvector and its derivative can be computed in a single algorithm.

Theorem 7.1. *If M is a matrix with only one simple eigenvalue of maximal modulus $\rho = |\lambda_1| > |\lambda_2|$, then the derivative g of the function $f \circ u$ at M , such that $g_{ij} = \sum_k \frac{\partial f}{\partial u_k} \frac{\partial u_k}{\partial M_{ij}}$ is the limit of the sequence $(\tilde{w}_l u_l^T)_{l \geq 0}$ given by the following iterative scheme:*

$$\begin{aligned} u_{l+1} &= \frac{Mu_l}{N(Mu_l)} \\ v_{l+1}^T &= \frac{v_l^T M}{v_l^T M u_{l+1}} \\ \tilde{w}_{l+1}^T &= \frac{1}{\rho_l} (\nabla f_l^T - (\nabla f_l \cdot u_l) \nabla N_l^T + \tilde{w}_l^T M)(I - u_{l+1} v_{l+1}^T) \end{aligned}$$

where $\rho_l = N(Mu_l)$, $\nabla f_l = \nabla f(u_l)$ and $\nabla N_l = \nabla N(u_l)$. Moreover, the sequences (u_l) , (v_l) and (\tilde{w}_l) converge linearly with rate $\frac{|\lambda_2|}{\rho}$.

Of course, the first and second sequences are the power method to the right and to the left. The third sequence is a modification the scheme of Proposition 7.9 with currently known values only. We shall denote one iteration of the scheme of the theorem as

$$(u_{k+1}, v_{k+1}, \tilde{w}_{k+1}) = \text{POWERDERIVATIVE}_M(u_k, v_k, \tilde{w}_k) .$$

Proof. The equalities giving u_{l+1} and v_{l+1} are simply the usual power method, so by [PPJ73], they convergence linearly with rate $\frac{|\lambda_2|}{\rho}$ to u and v , the right and left principal eigenvectors of M , such that $P = uv^T$ is the eigenprojector associated to $\rho(M)$. Let

$$z_l^T := \frac{1}{\rho_l} (-\nabla f_l^T + (\nabla f_l \cdot u_l) \nabla N_l^T) :$$

$\lim z_l = z$ by continuity of ∇f , N and ∇N at u . We also have

$$\tilde{w}_l^T = (-z_l^T + \frac{1}{\rho_l} \tilde{w}_l^T M)(I - u_{l+1} v_{l+1}^T) .$$

We first show that \tilde{w}_l is bounded. As in the proof of Proposition 7.9, $\rho(\tilde{M}(I - P)) = \frac{|\lambda_2|}{\rho} < 1$. Thus, by Lemma 5.6.10 in [HJ90], there exists a norm $\|\cdot\|_M$ and $\alpha < 1$ such that

$\tilde{M}(I - P)$ is α -contractant. Let S be the unit sphere: $\forall x \in S, \|\tilde{M}(I - P)x\|_M \leq \alpha\|x\|_M$. By continuity of the norm, $\forall \epsilon > 0, \exists L, \forall l \geq L, \forall x \in S, \|\frac{1}{\rho^l}M(I - u_{l+1}v_{l+1}^T)x\|_M \leq (\alpha + \epsilon)\|x\|_M$. As $\frac{1}{\rho^l}M(I - u_{l+1}v_{l+1}^T)$ is linear, we have the result on the whole \mathbb{R}^n space. Thus \tilde{w}_l is bounded.

Let us denote $\tilde{M} = \frac{1}{\rho}M$ and

$$\tilde{z}_l^T := z_l^T(I - u_{l+1}v_{l+1}^T) + \frac{1}{\rho^l}\tilde{w}_l^T M(uv^T - u_{l+1}v_{l+1}^T) + \frac{\rho - \rho^l}{\rho\rho^l}\tilde{w}_l^T M(I - uv^T),$$

so that $\tilde{w}_{l+1}^T = -\tilde{z}_l^T + \tilde{w}_l^T \tilde{M}(I - uv^T)$. We have:

$$\begin{aligned} \tilde{w}_l^T &= \tilde{w}_0^T (\tilde{M}(I - P))^l - \sum_{k=0}^{l-1} \tilde{z}_{l-1-k}^T (\tilde{M}(I - P))^k \\ &= \tilde{w}_0^T (\tilde{M}(I - P))^l - \sum_{k=0}^{l-1} z^T (\tilde{M}(I - P))^k - \sum_{k=0}^{l-1} (\tilde{z}_{l-1-k}^T - z^T) (\tilde{M}(I - P))^k \end{aligned}$$

By Proposition 7.9, the sum of the first and second summand correspond to w_l and converge linearly to w^T when l tends to infinity with convergence rate $\frac{|\lambda_2|}{\rho}$. Corollary 7.2 states that $g = \lim w_l u_l^T$. For the last one, we remark that $\forall \epsilon > 0, \|(\tilde{M}(I - P))^k\| = O(\frac{|\lambda_2 + \epsilon|}{\rho})^k$. In order to get the convergence rate of the sequence, we estimate $\|\tilde{z}_l - z\| = \|\tilde{z}_l - z_l + z_l - z\|$.

$$\|\tilde{z}_l - z\| \leq \|(z_l^T u_{l+1})v_{l+1}^T\| + \|\frac{1}{\rho^l}\tilde{w}_l^T M(uv^T - u_{l+1}v_{l+1}^T)\| + \|\frac{\rho - \rho^l}{\rho\rho^l}\tilde{w}_l^T\| + \|z_l - z\|$$

The second and third summands are clearly $O((\frac{|\lambda_2 + \epsilon|}{\rho})^l)$. For the first summand as $\nabla N_l^T u_l = 1$ we have

$$\begin{aligned} |z_l^T u_{l+1}| &= \left| \frac{1}{\rho^l} (-\nabla f_l^T u_{l+1} (\nabla f_l^T u_l) (\nabla N_l^T u_{l+1})) \right| \\ &\leq \frac{1}{\rho^l} (\sup_{l \geq 0} \|\nabla f_l^T\| + \sup_{l \geq 0} \|\nabla N_l^T\|) \|u_{l+1} - u_l\|. \end{aligned}$$

As $(u_l)_{l \geq 0}$ is bounded and f and N are C^1 , the constant is finite and $|z_l^T u_{l+1}| \|v_{l+1}\| = O((\frac{|\lambda_2 + \epsilon|}{\rho})^l)$. With similar arguments, we also show that $\|z_l - z\| = O((\frac{|\lambda_2 + \epsilon|}{\rho})^l)$.

Finally, we remark that for all $k, (\tilde{z}_{l-1-k}^T - z^T)(\tilde{M}(I - P))^k = O((\frac{|\lambda_2 + \epsilon|}{\rho})^{l-1})$. Thus

$$\sum_{k=0}^{l-1} (\tilde{z}_{l-1-k}^T - z^T)(\tilde{M}(I - P))^k = O\left(l \left(\frac{|\lambda_2 + \epsilon|}{\rho}\right)^{l-1}\right) = O\left(\left(\frac{|\lambda_2| + \epsilon'}{\rho}\right)^l\right),$$

for all $\epsilon' > \epsilon$. The result follows. \square

Remark 7.3. One iteration of the algorithm presented in Theorem 7.1 is composed of 3 multiplications by the matrix M (assumed sparse with filling rate $\gamma(M)$), 1 gradient evaluation of the utility (cost $c(f)$), 1 evaluation of the normalization function and of its gradient (cost $c(N)$), 3 scalar products, 2 vector sums and 3 scalar-vector multiplications. Hence, the complexity is $3\gamma(M)n^2 + 8n + c(f) + c(N)$. The leading term is the sparse matrix-vector multiplication. This is the same order of complexity as for the power method, which is $\gamma(M)n^2 + n + c(N)$. Moreover, as the rate of convergence of both algorithms is equal, the number of iterations is comparable.

Remark 7.4. All this applies easily to a nonnegative irreducible matrix M . Let ρ be its principal eigenvalue: it is simple thanks to irreducibility. The spectral gap assumption $\rho > |\lambda_2|$ is guaranteed by an additional aperiodicity assumption. Let v and u be the left and right eigenvectors of M for the eigenvalue ρ . We normalize u by $N(u) = 1$ where N verifies $\frac{\partial N}{\partial u}(u) \geq 0$ and $N(\lambda u) = \lambda N(u)$ (which implies $\frac{\partial N}{\partial u}(u) \cdot u = N(u)$) and v by $\sum_i u_i v_i = 1$. As $u > 0$, any normalization such that $p = \frac{\partial N}{\partial u}(u) \geq 0$ is satisfactory: for instance, we could choose $N(u) = \|u\|_1 = \sum_i u_i$, $N(u) = \|u\|_2$ or $N(u) = u_1$.

Remark 7.5. Theorem 7.1 gives the possibility of performing a gradient algorithm for Perron vector optimization. Fix a precision $\epsilon > 0$ and apply recursively the power-type iterations POWERDERIVATIVE_M until $\|u_l - u_{l+1}\| + \|\tilde{w}_l - \tilde{w}_{l+1}\| \leq \epsilon$. Then we use $\tilde{w}_l u_l$ as the descent direction of the algorithm. The gradient algorithm will stop at a nearly stationary point, the smaller ϵ the better. In order to accelerate the algorithm, we can initialize the recurrence with former values of u_l , v_l and \tilde{w}_l .

7.4 Coupling gradient and power iterations

We have given in Theorem 7.1 an algorithm that gives the derivative of the objective function with the same order of complexity as the computation of the value of the function by the power method. As the problem is a differentiable optimization problem, we can perform any classical optimization algorithm: see [BGLS06, Ber95, NW99] for references.

When we consider relaxations of the HITS authority optimization problem, that we will define in Sections 7.5 (or another Perron ranking), the constraints on the adjacency matrices are very easy to deal with, so a projected gradient algorithm as described in [Ber76] will be efficient. If the problem has not a too big size, it is also possible to set a second order algorithm. However, matrices arising from web applications are large: as the Hessian matrix is a full $n^2 \times n^2$ matrix, it is then difficult to work with.

In usual algorithms, the value of the objective function must be evaluated at any step of the algorithm. As stressed in [Ove91], there are various possibilities for the computation of the eigenvalue and eigenvectors. Here, we consider sparse nonnegative matrices with a simple principal eigenvalue: the power method applies and, unlike direct methods or inverse iterations, it only needs matrix-vector products, which is valuable with a large sparse matrix. Nevertheless for large matrices, repeated principal eigenvector and eigenvalue determinations can be costly. Hence, we give a first order descent algorithm designed to find stationary points of Perron eigenvector optimization problems, that uses approximations of the value of the objective and of its gradient instead of the precise values. Then we can interrupt the computation of the eigenvector and eigenvalue when necessary and avoid useless computations. Moreover, as the objective is evaluated as the limit of a sequence, its exact value is not available in the present context.

The algorithm consists of a coupling of the power iterations and of the gradient algorithm with Armijo line search along the projected arc [Ber76]. We recall this latter gradient algorithm in Algorithm 7.1. We shall, instead of comparing the exact values of the function, compare approximations of these values, the level of approximation being monitored during the course of the algorithm.

If we had an easy access to the exact value of $u(M)$ for all $M \in h(\mathcal{C})$, we could use the gradient algorithm with Armijo line search along the projected arc with $J = f \circ u \circ h$ to find a stationary point of Problem (7.2). But here, we compute the Perron eigenvector and its

Algorithm 7.1 Gradient algorithm with Armijo line search along the projected arc [Ber76] Let a differentiable function J , a convex admissible set \mathcal{C} and an initial point $x_0 \in \mathcal{C}$ and parameters $\sigma \in (0, 1)$, $\alpha^0 > 0$ and $\beta \in (0, 1)$. The algorithm is an iterative algorithm defined for all $k \in \mathbb{N}$ by

$$x_{k+1} = \text{P}_{\mathcal{C}}(x_k - \alpha_k \nabla J(x_k))$$

and $\alpha_k = \beta^{m_k} \alpha^0$ where m_k is the first nonnegative integer m such that

$$J(\text{P}_{\mathcal{C}}(x_k - \beta^m \alpha^0 \nabla J(x_k))) - J(x_k) \leq -\sigma \frac{\|x_k - \text{P}_{\mathcal{C}}(x_k - \beta^m \alpha^0 \nabla J(x_k))\|_2^2}{\beta^m \alpha^0} .$$

derivatives by an iterative scheme (Theorem 7.1). So we only have converging approximations of the value of the objective and of its gradient and we cannot use directly the classical gradient algorithm (the approximate gradient may not be a descent direction). The theory of consistent approximation, developed in [Pol97] proposes algorithms and convergence results for such problems. If the main applications of consistent approximations are optimal control and optimal control of partial derivative equations, it is also useful for problems in finite dimension where the objective is difficult to compute [PP02].

A consistent approximation of a given optimization problem is a sequence of computationally tractable problems that converge to the initial problem in the sense that the stationary points of the approximate problems converge to stationary points of the original problem. The theory provides master algorithms that construct a consistent approximation, initialize a nonlinear programming algorithm on this approximation and terminate its operation when a precision (or discretization) improvement test is satisfied.

We consider the Perron vector optimization problem defined in (7.2)

$$\min_{x \in \mathcal{C}} J(x) = \min_{x \in \mathcal{C}} f \circ u \circ h(x) .$$

For $x \in \mathcal{C}$, $n \in \mathbb{N}$ and for arbitrary fixed vectors u_0 , v_0 and \tilde{w}_0 , we shall approximate with order $\Delta(n)$ the Perron vectors of $h(x)$, namely u and v and the auxiliary vector w of Corollary 7.2 by

$$(u_{k_n}, v_{k_n}, \tilde{w}_{k_n}) := (\text{POWERDERIVATIVE}_{h(x)})^{k_n}(u_0, v_0, \tilde{w}_0) , \quad (7.6)$$

where k_n is the first nonnegative integer k such that

$$\|(u_{k+1}, v_{k+1}, \tilde{w}_{k+1}) - (u_k, v_k, \tilde{w}_k)\| \leq \Delta(n) . \quad (7.7)$$

The map POWERDERIVATIVE is defined in Theorem 7.1. Then the degree n approximation of the objective function J and of its gradient ∇J are given by

$$J_n(x) = f(u_{k_n}) , \quad g_n(x) = \sum_{i,j \in [n]} \tilde{w}_{k_n}(i) \nabla h_{i,j}(x) u_{k_n}(j) . \quad (7.8)$$

An alternative approach, proposed in [PP02], is to approximate (u, v, w) by the n^{th} iterate $(u_n, v_n, \tilde{w}_n) := (\text{POWERDERIVATIVE}_{h(x)})^n(u_0, v_0, \tilde{w}_0)$. We did not choose this approach since it does not take into account efficiently hot started power iterations.

The usual Armijo line search requires the exact value of the gradient in order to terminate in a finite number of iterations. Here, as we have only approximations of this gradient, the

classical Armijo rule may fail. Hence we define the step with an interrupted Armijo line search in Algorithm 7.2. The main difference with the usual Armijo line search is that, depending on the current precision parameter n , we define a maximum number of step decreases \bar{M}_n . If this number of decreases is not sufficient to satisfy Armijo conditions, then we declare that the line search has failed and we force the master algorithm to increase the precision.

Algorithm 7.2 Interrupted Armijo line search along the projected arc

Let $(\bar{M}_n)_{n \geq 0}$ be a sequence diverging to $+\infty$, $\sigma \in (0, 1)$, $\alpha^0 > 0$, $\beta \in (0, 1)$ and $\gamma > 0$. Given $n \in \mathbb{N}$, J_n and g_n are defined in (7.8). For $x \in \mathcal{C}$, the algorithm returns $A_n(x)$ defined as follows. If for all nonnegative integer m smaller than \bar{M}_n ,

$$J_n(\text{P}_{\mathcal{C}}(x - \beta^m \alpha^0 g_n(x))) - J_n(x) > -\sigma \frac{\|x - \text{P}_{\mathcal{C}}(x - \beta^m \alpha^0 g_n(x))\|_2^2}{\beta^m \alpha^0}$$

then we say that the line search has failed and we set $A_n(x) = \emptyset$. Otherwise, let m_n be the first nonnegative integer m such that

$$J_n(\text{P}_{\mathcal{C}}(x - \beta^{m_n} \alpha^0 g_n(x))) - J_n(x) \leq -\sigma \frac{\|x - \text{P}_{\mathcal{C}}(x - \beta^{m_n} \alpha^0 g_n(x))\|_2^2}{\beta^{m_n} \alpha^0}$$

and define the next iterate $A_n(x)$ to be $A_n(x) = \text{P}_{\mathcal{C}}(x - \beta^{m_n} \alpha^0 g_n(x))$.

Then we shall use the Interrupted Armijo line search along the projected arc A_n in the following Master Algorithm Model (Algorithm 7.3), where the set of precisions for which the Interrupted Armijo line search does not fail is denoted by the set \mathcal{N} .

Algorithm 7.3 Master Algorithm Model 3.3.17 in [Pol97]

Let $\omega \in (0, 1)$, $\sigma' \in (0, 1)$, $n_{-1} \in \mathbb{N}$ and $x_0 \in \mathcal{C}$, $\mathcal{N} = \{n \mid A_n(x) \neq \emptyset, \forall x \in \mathcal{C}\}$ and $(\Delta(n))_{n \geq 0}$ be a sequence converging to 0.

For $i \in \mathbb{N}$, compute iteratively the smallest $n_i \in \mathcal{N}$ and x_{i+1} such that $n_i \geq n_{i-1}$,

$$\begin{aligned} x_{i+1} &\in A_{n_i}(x_i) \quad \text{and} \\ J_{n_i}(x_{i+1}) - J_{n_i}(x_i) &\leq -\sigma' \Delta(n_i)^\omega. \end{aligned}$$

In order to prove the convergence of the Master Algorithm Model (algorithm 7.3) when used with the Interrupted Armijo line search (Algorithm 7.2), we need the following lemma.

Lemma 7.2. *For all $x^* \in \mathcal{C}$ which is not stationary, there exists $\rho^* > 0$, $\delta^* > 0$ and $n^* \in \mathbb{N}$ such that for all $n \geq n^*$, and for all $x \in \text{B}(x^*, \rho^*) \cap \mathcal{C}$, $A_n(x) \neq \emptyset$ and*

$$J_n(\text{P}_{\mathcal{C}}(x - \alpha_n g_n(x))) - J_n(x) \leq -\delta^*$$

where α_n is the step length returned by the Interrupted Armijo line search $A_n(x)$ (Algorithm 7.2).

Proof. Let $x \in \mathcal{C}$. Suppose that there exists an infinitely growing sequence $(\phi_n)_{n \geq 0}$ such that $A_{\phi_n}(x) = \emptyset$ for all n . Then for all $m \leq \bar{M}_{\phi_n}$,

$$J_{\phi_n}(\text{P}_{\mathcal{C}}(x - \beta^m \alpha^0 g_{\phi_n}(x))) - J_{\phi_n}(x) > -\sigma \langle g_{\phi_n}(x), x - \text{P}_{\mathcal{C}}(x - \beta^m \alpha^0 g_{\phi_n}(x)) \rangle$$

When $n \rightarrow +\infty$, $\bar{M}_{\phi_n} \rightarrow +\infty$, $J_{\phi_n}(x) \rightarrow J(x)$ and $g_{\phi_n}(x) \rightarrow \nabla f(x)$ (Theorem 7.1), so we get that for all $m \in \mathbb{N}$, $J(\mathcal{P}_{\mathcal{C}}(x_k - \beta^m \alpha^0 \nabla J(x_k))) - J(x_k) \geq -\sigma \frac{\|x_k - \mathcal{P}_{\mathcal{C}}(x_k - \beta^m \alpha^0 \nabla J(x_k))\|_2^2}{\beta^m \alpha^0}$, which is impossible by [Ber76].

So suppose $n \in \mathbb{N}$ is sufficiently large so that $A_n(x) \neq \emptyset$. Let α_n the step length determined by Algorithm 7.2 and let α be the step length determined by Algorithm 7.1 at x . We have:

$$J_n(\mathcal{P}_{\mathcal{C}}(x - \alpha_n g_n(x))) - J_n(x) \leq -\sigma \frac{\|x - \mathcal{P}_{\mathcal{C}}(x - \alpha_n g_n(x))\|_2^2}{\alpha_n}$$

and if $\alpha_n \neq \alpha^0$,

$$J_n(\mathcal{P}_{\mathcal{C}}(x - \beta^{-1} \alpha_n g_n(x))) - J_n(x) > -\sigma \frac{\|x - \mathcal{P}_{\mathcal{C}}(x - \beta^{-1} \alpha_n g_n(x))\|_2^2}{\beta^{-1} \alpha_n}.$$

$(\alpha_n)_{n \geq 0}$ is a bounded sequence so it has a subsequence $(\alpha_{\phi_n})_{n \geq 0}$ converging to, say, $\bar{\alpha}$. As $(\alpha_n)_{n \geq 0}$ can only take discrete values, this means that $\alpha_{\phi_n} = \bar{\alpha}$ for all n sufficiently big.

When n tend to infinity, by Theorem 7.1, we get

$$J(\mathcal{P}_{\mathcal{C}}(x - \bar{\alpha} \nabla f(x))) - J(x) \leq -\sigma \frac{\|x - \mathcal{P}_{\mathcal{C}}(x - \bar{\alpha} \nabla f(x))\|_2^2}{\bar{\alpha}}$$

and if $\bar{\alpha} \neq \alpha^0$,

$$J(\mathcal{P}_{\mathcal{C}}(x - \beta^{-1} \bar{\alpha} \nabla f(x))) - J(x) \geq -\sigma \frac{\|x - \mathcal{P}_{\mathcal{C}}(x - \beta^{-1} \bar{\alpha} \nabla f(x))\|_2^2}{\beta^{-1} \bar{\alpha}}.$$

Then, if α is the step length returned by Armijo rule (Algorithm 7.1), then $\alpha \geq \bar{\alpha}$, because α is the first number of the sequence that verifies the first inequality. Similarly, consider the version of Armijo rule with a strict inequality instead of the non strict inequality. Then if α_{strict} is the step length returned by this algorithm, we have $\alpha_{\text{strict}} \leq \bar{\alpha}$.

Moreover, like in [PP02] one can easily see that $\forall x^* \in \mathcal{C}$ not stationary, $\exists \rho^* > 0$, $\exists \delta^* > 0$, such that $\forall x \in \mathcal{B}(x^*, \rho^*) \cap \mathcal{C}$,

$$J((\mathcal{P}_{\mathcal{C}}(x - \alpha_{\text{strict}} \nabla f(x)))) - J(x) \leq -\delta^*$$

where $\delta^* = \sigma \frac{\|x^* - \mathcal{P}_{\mathcal{C}}(x^* - \alpha_{\text{strict}} \nabla f(x^*))\|_2^2}{\alpha_{\text{strict}}} - \|\nabla J(x^*)\| \rho^*$. By Lemma 3 in [GB82], $\alpha_{\text{strict}} \leq \bar{\alpha}$ implies that $\frac{\|x^* - \mathcal{P}_{\mathcal{C}}(x^* - \alpha_{\text{strict}} \nabla f(x^*))\|_2^2}{\alpha_{\text{strict}}} \geq \frac{\|x - \mathcal{P}_{\mathcal{C}}(x - \bar{\alpha} \nabla f(x))\|_2^2}{\bar{\alpha}}$. As this is true for all adherent point of $(\alpha_n)_{n \geq 0}$, $\forall x^* \in \mathcal{C}$ not stationary, $\exists \rho^* > 0$, $\exists \delta^* > 0$ and $\exists n^* \in \mathbb{N}$, such that $\forall n \geq n^*$, $\forall x \in \mathcal{B}(x^*, \rho^*) \cap \mathcal{C}$,

$$J_n((\mathcal{P}_{\mathcal{C}}(x - \alpha_n g_n(x)))) - J_n(x) \leq -\sigma \frac{\|x - \mathcal{P}_{\mathcal{C}}(x - \alpha_{\text{strict}} \nabla f(x))\|_2^2}{\alpha_{\text{strict}}} + \delta^*/4 \leq -\delta^*/2,$$

for n^* sufficiently large and ρ^* sufficiently small. \square

In [PP02], the property of the lemma was proved for exact minimization in the line search. We proved it for Algorithm 7.2.

We shall also need the following result

Proposition 7.10 (Theorem 25 in [May94]). *Let $M \in M_{n,n}(\mathbb{R})$, $\tilde{\lambda} \in \mathbb{R}$, $\tilde{x} \in \mathbb{R}^n$, $C \in M_{n+1,n+1}(\mathbb{R})$ and $p \in \mathbb{R}^n$ such that $p^T \tilde{x} = 1$. Denote*

$$B = \begin{bmatrix} A - \tilde{\lambda}I_n & -\tilde{x} \\ p^T & 0 \end{bmatrix}, \quad \eta = \left\| C \begin{bmatrix} A\tilde{x} - \tilde{\lambda}\tilde{x} \\ 0 \end{bmatrix} \right\|_{\infty},$$

$\sigma = \|I_{n+1} - CB\|$ and $\tau = \|C\|_{\infty}$. *If $\sigma < 1$ and $\Delta = (1 - \sigma)^2 - 4\eta\tau \geq 0$, then $\beta = \frac{2\eta}{1 - \sigma + \sqrt{\Delta}}$ is nonnegative and there exists a unique eigenpair (x^*, λ^*) of M such that $p^T x^* = 1$, $|\lambda^* - \tilde{\lambda}| \leq \beta$ and $\|x^* - \tilde{x}\|_{\infty} \leq \beta$.*

Theorem 7.2. *Let $(x_i)_{i \geq 0}$ be a sequence constructed by the Master Algorithm Model (Algorithm 7.3) for the resolution of the Perron vector optimization problem (7.2) such that $A_n(x)$ is the Interrupted Armijo line search along the projected arc (Algorithm 7.2) and $\Delta(n) = (\Delta_0)^n$ for $\Delta_0 \in (0, 1)$. Then every accumulation point of $(x_i)_{i \geq 0}$ is a stationary point of (7.2).*

Proof. The proof of the theorem is based on Theorem 3.3.19 in [Pol97]. This theorem shows that if continuity assumptions hold (they trivially hold in our case), if for all bounded subset S of \mathcal{C} there exist $K > 0$ such that for all $x \in \mathcal{C}$

$$|J(x) - J_n(x)| \leq K\Delta(n),$$

and if for all $x^* \in \mathcal{C}$ which is not stationary, there exists $\rho^* > 0$, $\delta^* > 0$ and $n^* \in \mathbb{N}$ such that for all $n \geq n^*$, for all $x \in B(x^*, \rho^*) \cap \mathcal{C}$ and for all $y \in A_n(x)$,

$$J_n(y) - J_n(x) \leq -\delta^*,$$

then every accumulation point of a sequence $(x_i)_{i \geq 0}$ generated by the Master Algorithm Model (algorithm 7.3) is a stationary point of the problem of minimizing $J(x)$.

We first remark that for $x \in \mathcal{C}$, $u = u(h(x))$ and J_n defined in (7.8), as f is continuously differentiable, we have $|J(x) - J_n(x)| \leq \|\nabla f(u)\| \|u - u_{k_n}\|$.

We shall now show that for all matrix M , there exists $K > 0$ such that $\|u - u_n\| \leq K\|u_{n+1} - u_n\|$. Remark that Theorem 16 in [May94] gives the result with $K > \frac{N(u_n)}{\rho - |\lambda_2|}$ when M is symmetric. When M is not necessarily symmetric, we use Proposition 7.10 with $\tilde{x} = u$, $\tilde{\lambda} = \rho$ and C is the inverse of B . Let $\epsilon > 0$, by continuity of the inverse, for n sufficiently large, if we define B_n to be the matrix of Proposition 7.10 with $\tilde{x} = u_n$ and $\tilde{\lambda} = N(Mu_n)$, we still have $\sigma := \|I_{n+1} - CB_n\| < \epsilon$. We also have:

$$\begin{aligned} \eta &:= \left\| C \begin{bmatrix} M\tilde{x} - \tilde{\lambda}\tilde{x} \\ 0 \end{bmatrix} \right\|_{\infty} \leq \|C\|_{\infty} \|Mu_n - N(Mu_n)u_n\|_{\infty} \\ &\leq \|C\|_{\infty} N(Mu_n) \|u_{n+1} - u_n\|. \end{aligned}$$

The conclusion of Proposition 7.10 tells us that if $\Delta := (1 - \sigma)^2 - 4\eta\tau \geq 0$, then

$$\|u - u_n\| \leq \beta := \frac{2\eta}{1 - \sigma + \sqrt{\Delta}} \leq \frac{2\eta}{1 - \sigma} \leq 3\eta \leq 3\|C\|_{\infty} N(Mu_n) \|u_{n+1} - u_n\|.$$

Now, as the inversion is a continuous operation, for all compact subset S of \mathcal{C} , there exists $K > 0$ such that for all $M = h(x) \in h(\mathcal{C})$ and for all n sufficiently big, $\|u - u_n\| \leq K\|u_{n+1} - u_n\|$.

By definition of k_n (7.7), we get

$$|J(x) - J_n(x)| \leq \|\nabla f(u)\| \|u - u_{k_n}\| \leq K\|u_{k_n+1} - u_{k_n}\| \leq K\Delta(n).$$

By Lemma 7.2, the other hypothesis of Theorem 3.3.19 in [Pol97] is verified and the conclusion holds: every accumulation point of $(x_i)_{i \geq 0}$ is a stationary point of (7.2). \square

7.5 Application to HITS optimization

In the last two sections, we have developed scalable algorithms for the computation of the derivative of a scalar function of the Perron vector of a matrix and for the searching of stationary points of Perron vector optimization problems (7.2). We now apply these results to two web ranking optimization problems, namely HITS authority optimization and HOTS optimization.

HITS algorithm for ranking web pages has been described by Kleinberg in [Kle99]. The algorithm has two phases: first, given a query, it produces a subgraph G of the whole web graph such that it contains relevant pages, pages linked to relevant pages and the hyperlinks between them. The second phase consists in computing a principal eigenvector called authority vector and to sort the pages with respect to their corresponding entry in the eigenvector. If we denote by A the adjacency matrix of the directed graph G , then the authority vector is the principal eigenvector of $A^T A$.

It may however happen that $A^T A$ is reducible and then the authority vector is not uniquely defined. Following [LM06], we remedy this by defining the HITS authority score to be the principal eigenvector of $A^T A + \xi ee^T$, for a given small positive real ξ . We then normalize the HITS vector with the 2-norm as proposed by Kleinberg [Kle99].

Given a subgraph associated to a query, we study in this section the optimization of the authority of a set of pages. We partition the set of potential links (i, j) into three subsets, consisting respectively of the set of obligatory links \mathcal{O} , the set of prohibited links \mathcal{I} and the set of facultative links \mathcal{F} . Some authors consider that links between pages of a website, called intra-links, should not be considered in the computation of HITS. This results in considering these links as prohibited because this is as if they did not exist.

Then, we must select the subset J of the set of facultative links \mathcal{F} which are effectively included in this page. Once this choice is made for every page, we get a new webgraph, and define the adjacency matrix $A = A(J)$. We make the simplifying assumption that the construction of the focused graph G is independent of the set of facultative links chosen.

Given a utility function f , the HITS authority optimization problem is:

$$\max_{J \subseteq \mathcal{F}, u \in \mathbb{R}^n, \lambda \in \mathbb{R}} \{f(u) ; (A(J)^T A(J) + \xi ee^T)u = \lambda u, \|u\|_2 = 1, u \geq 0\} \quad (7.9)$$

The set of admissible adjacency matrices is a combinatorial set with a number of matrices exponential in the number of facultative links. Thus we shall consider instead a relaxed version of the HITS authority optimization problem which consists in accepting weighted adjacency matrices. It can be written as

$$\begin{aligned} & \max_{A \in \mathbb{R}^{n \times n}, u \in \mathbb{R}^n, \rho \in \mathbb{R}} f(u) \\ & (A^T A + \xi ee^T)u = \rho u, \quad \|u\|_2 = 1, \quad u \geq 0 \\ & A_{i,j} = 1, \quad \forall (i, j) \in \mathcal{O} \\ & A_{i,j} = 0, \quad \forall (i, j) \in \mathcal{I} \\ & 0 \leq A_{i,j} \leq 1, \quad \forall (i, j) \in \mathcal{F} \end{aligned} \quad (7.10)$$

The relaxed HITS authority optimization problem (7.10) is a Perron vector optimization problem (7.2) with $h(A) = A^T A + \xi ee^T$ and the normalization $N(u) = \sqrt{\sum_i u_i^2} = 1$. Hence $\nabla N(u(M)) = u(M)$. Remark that $\|u\| = \|v\| = 1$. Now $\frac{\partial h}{\partial A}(A).H = H^T A + A^T H$ so the

derivative of the criterion with respect to the weighted adjacency matrix is $(Aw)u^T + (Au)w^T$ with $w = (\nabla f^T - (\nabla f \cdot u)\nabla N^T)(A^T A + \xi ee^T - \rho I)^\#$.

Thanks to $\xi > 0$, the matrix is irreducible and aperiodic. Thus, it has only one eigenvalue of maximal modulus and we can apply Theorem 7.2.

The next proposition shows that, as is the case for PageRank optimization [dKNvD08, FABG13], optimal strategies have a rather simple structure.

Proposition 7.11 (Threshold property). *Let A be a locally maximal linking strategy of the relaxed HITS authority optimization problem (7.10) with associated authority vector u and derivative at optimum $(Aw)u^T + (Au)w^T$. For all controlled page i denote $b_i = \frac{-(Aw)_i}{(Au)_i}$ if it has at least one outlink. Then all facultative hyperlinks (i, j) such that $\frac{w_j}{u_j} > b_i$ get a weight of 1 and all facultative hyperlinks (i, j) such that $\frac{w_j}{u_j} < b_i$ get a weight of 0.*

In particular, if two pages with different b_i 's have the same sets of facultative outlinks, then their set of activated outlinks are included one in the other.

Proof. As the problem only has box constraints, a nonzero value of the derivative at the maximum determines whether the upper bound is saturated ($g_{i,j} > 0$) or the lower bound is saturated ($g_{i,j} < 0$). If the derivative is zero, the weight of the link can take any value.

We have $g_{i,j} = (Aw)_i u_j + (Au)_i w_j$ with $u_j > 0$ and $(Au)_i \geq 0$. If Page i has at least one outlink, then $(Au)_i > 0$ and we simply divide by $(Au)_i u_j$ to get the result thanks to the first part of the proof. If two pages i_1 and i_2 have the same sets of facultative outlinks and if $b_{i_1} < b_{i_2}$, then $\frac{w_j}{u_j} \geq b_{i_2}$ implies $\frac{w_j}{u_j} > b_{i_1}$: all the pages pointed by i_2 are also pointed by i_1 . \square

Remark 7.6. If a page i has no outlink, then $(Aw)_i = (Au)_i = 0$ and $g_{i,j} = 0$ for all $j \in [n]$, so we cannot conclude with the argument of the proof.

Remark 7.7. This proposition shows that $\frac{w_j}{u_j}$ gives a total order of preference in pointing to a page or another.

Then we give on Figures 7.1 and 7.2 a simple HITS authority optimization problem and two local solutions. They show the following properties for this problem.

Example 7.1. The relaxed HITS authority optimization problem is in general not quasi-convex nor quasi-concave.

Proof. Any strict local maximum of a quasi-convex problem is necessarily an extreme of the admissible polyhedral set (this is a simple extension of Theorem 3.5.3 in [BS06]). The example on Figure 7.1 shows that this is not the case here.

A quasi-concave problem can have only one strict local maximum (although it may have many local maxima). The examples on Figures 7.1 and 7.2 show two distinct strict local maxima for a HITS authority optimization problem. \square

Heuristic 7.1. These examples also show that the relaxed HITS authority optimization problem (7.10) does not give binary solutions that would be then solutions of the initial HITS authority optimization problem (7.9). Hence we propose the following heuristic to get “good” binary solutions. From a stationary point of the relaxed problem, define the function $\phi : [0, 1] \rightarrow \mathbb{R}$ such that $\phi(x)$ is the value of the objective function when we select in (7.9) all the links with weight greater than x in the stationary point of (7.10). We only need to compute it at a finite number of points since ϕ is piecewise constant. We then select the best threshold.




Figure 7.1: Strict local maximum for relaxed HITS authority optimization on a small web graph of 21 pages with 3 controlled pages (colored) representing the website I . Obligatory links are the thin arcs, facultative links are all the other outlinks from the controlled pages except self links. The locally optimal solution for the maximization of $f(u) = \sum_{i \in I} u_i^2$ is to select the bold arcs with weight 1 and the dotted arc with weight 0.18. Selected internal links are dark blue, selected external links are light red. We checked numerically the second order optimality conditions [Ber95]. The initial sum of HITS authority scores was 0.056. The best strategy we found gives a score of 0.357 (the maximum possible is 1). This large increase lets the controlled pages rank 1st, 4th and 9th over 21 in terms of HITS authority instead of 4th, 9th and 19th initially.




Figure 7.2: Another strict local maximum for the same HITS authority optimization problem as in Figure 7.1.

For instance, with the stationary point of Figure 7.1, this heuristic suggests not to select the weighted link.

7.6 Optimization of the eigenvector of a nonlinear Perron-Frobenius operator

Some web ranking algorithms, like sinkhorning [Smi05, Kni08] or the ideal HOTS algorithm [Tom03] are defined as the eigenvector of a nonlinear Perron-Frobenius operator. In this section, we show that the main properties of linear Perron vector optimization, i.e. the low rank of the matrix of partial derivatives and the sparse calculus, can be adapted to the nonlinear case.

We recall the following simple result:

Lemma 7.3. *Let T be a monotone homogeneous map. Let u and λ be such that $T(u) = \lambda u$. Then the matrix $\frac{\partial T}{\partial u}(u)$ is nonnegative, its Perron eigenvalue is λ and u is one of its Perron eigenvector.*

Proof. As for all $t > 0$, $u + te_j \geq u$ and T is monotone,

$$\frac{\partial T_i}{\partial u_j}(u) = \lim_{t \rightarrow 0, t > 0} \frac{T_i(u + te_j) - T_i(u)}{t} \geq 0 .$$

For $x \in \mathbb{R}$, we differentiate $T(xu) = x\lambda u$ with respect to x to get:

$$\frac{\partial T}{\partial u}(xu)u = \lambda u .$$

When we evaluate this equality at $x = 1$, we deduce that (λ, u) is an eigenpair of $\frac{\partial T}{\partial u}(u)$. As λ and u are nonnegative, this means that λ is the Perron eigenvalue of the matrix. \square

Definition 7.2. We shall say that a family $(T(\cdot, x))_x$ is a differentiable family of Perron-Frobenius operators if $((u, x) \mapsto T(u, x))$ is continuously differentiable, if for all x , the function $(u \mapsto T(u, x))$ is monotone, homogeneous and if $\frac{\partial T}{\partial u}(u, x)$ is primitive at a given eigenvector u of $T(\cdot, x)$. This last assumption implies that $T(\cdot, x)$ has a unique eigenvector that we shall denote $u(x)$.

The nonlinear Perron-Frobenius eigenvector optimization problem can be written as

$$\min_{x \in \mathcal{C}} f(u(x)) \quad (7.11)$$

We assume that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a real valued continuously differentiable function and that \mathcal{C} a closed convex set of parameters such that the map $((u, x) \mapsto T(u, x))$ is continuously differentiable on $\mathbb{R}^n \times \mathcal{C}$.

Proposition 7.12. *Let $(T(\cdot, x'))_{x'}$ be a differentiable family of Perron-Frobenius operators. Given a parameter x , let λ and u be the eigenvalue and eigenvector of $T(\cdot, x)$ such that u is normalized by $N(u) = 1$. We denote $S = (\frac{\partial T}{\partial u}(u, x) - \lambda I)^\#$. Then the partial derivatives of the eigenvector are given by:*

$$\frac{\partial u}{\partial x_k}(x) = -S \frac{\partial T}{\partial x_k}(u, x) + (\nabla N(u))^T S \frac{\partial T}{\partial x_k}(u, x) u$$

Proof. We differentiate the equality $T(u, x) = \lambda u$ with respect to x_k

$$\frac{\partial T}{\partial u} \frac{\partial u}{\partial x_k} + \frac{\partial T}{\partial x_k} = \frac{\partial \lambda}{\partial x_k} u + \lambda \frac{\partial u}{\partial x_k}$$

Let v be the normalized left Perron vector of $\frac{\partial T}{\partial u}$ such that $S(\frac{\partial T}{\partial u} - \lambda I) = (I - uv^T)$. We have, using the fact that $Su = 0$,

$$(I - uv^T) \frac{\partial u}{\partial x_k} = -S \frac{\partial T}{\partial x_k} \quad (7.12)$$

As $\nabla N^T \frac{\partial u}{\partial x_k} = 0$ and $\nabla N^T u = 1$, we get

$$-v^T \frac{\partial u}{\partial x_k} = -\nabla N^T S \frac{\partial T}{\partial x_k}$$

We reinject this last equality in (7.12) to get

$$\frac{\partial u}{\partial x_k} = -S \frac{\partial T}{\partial x_k} + (\nabla N^T S \frac{\partial T}{\partial x_k}) u \quad \square$$

To simplify notations, we denote ∇f for $\nabla f(u)$, ∇N for $\nabla N(u)$ and $\frac{\partial T}{\partial x}$ for $\frac{\partial T}{\partial x}(u(x), x)$.

Corollary 7.3. *With the notations of Proposition 7.12, the partial derivatives of the function $x \mapsto (f \circ u)(x)$ at x are given by $g_k = \sum_{i \in [n]} w_i \frac{\partial T_i}{\partial x_k}$, where the auxiliary vector w is given by:*

$$w^T = (-\nabla f^T + (\nabla f \cdot u) \nabla N^T) S = (-\nabla f^T + (\nabla f \cdot u) p^T) (\frac{\partial T}{\partial u} - \lambda I)^\#$$

Proof. Similar to the proof of Corollary 7.2. \square

Since $\frac{\partial T}{\partial u}$ is a primitive matrix, the auxiliary vector can be computed with the iterative scheme of Proposition 7.9. However, without any additional information on the structure of $\frac{\partial T}{\partial x}$, we cannot simplify the result further. If x is a matrix of size n , we may thus get a large matrix of partial derivatives. We shall then prove that this matrix has a small rank to make the algorithm scalable. This is the case for sinkhorning or the ideal HOTS algorithm for instance.

In the next section, we adapt these results to the optimization of the HOTS score. Then, $\frac{\partial T}{\partial x}$ is a matrix of rank at most 3 and thus Proposition 7.3 gives an efficient formula. In fact, the effective HOTS operator is not monotone but the main results useful for the optimization of the ranking still hold.

7.7 Optimization of the HOTS score of a website's pages

In this section, we study the optimization of the HOTS score of a website's pages. Then the fixed point operator is not monotone but we show that the matrix of partial derivatives has rank at most 3 and that the Master Algorithm model can be used for this problem. We think that this supports Tomlin's remark that "malicious manipulation of the dual values of a large scale nonlinear network optimization model [...] would be an interesting topic".

As for HITS authority in Section 7.5, we consider sets of obligatory links, prohibited links and facultative links. From now on, the adjacency matrix A may change, so we define $\tilde{\theta}$ and F as functions of p and A . For all A , the HOTS vector is uniquely defined up to an additive constant for $\alpha < 1$, so we shall set a normalization, like for instance $\sum_i p_i = 0$ or $\log(\sum_i \exp(p_i)) = 0$. Thus, given a normalization function N , we can define the function $p : A \mapsto p(A)$. For all A, i, j , the normalization function N may verify $\frac{\partial N}{\partial p}(p(A)) \frac{\partial p}{\partial A_{i,j}}(A) = 0$ and $N(p + \lambda) = N(p) + \lambda$ for all $\lambda \in \mathbb{R}$, so that $\frac{\partial N}{\partial p}(p)e = 1$.

The HOTS authority optimization problem is:

$$\max_{J \subseteq \mathcal{F}, p \in \mathbb{R}^n, \lambda \in \mathbb{R}} \{f(p) ; u(A(J), p) = p, N(p) = 0, \} \quad (7.13)$$

We shall mainly study instead the relaxed HOTS optimization problem which can be written as:

$$\begin{aligned} & \max_{A \in \mathbb{R}^{n \times n}, p \in \mathbb{R}^n} f(u) \\ & u(A, p) = p, \quad N(p) = 0 \\ & A_{i,j} = 1, \quad \forall (i, j) \in \mathcal{O} \\ & A_{i,j} = 0, \quad \forall (i, j) \in \mathcal{I} \\ & 0 \leq A_{i,j} \leq 1, \quad \forall (i, j) \in \mathcal{F} \end{aligned} \quad (7.14)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is the objective function. We will assume that f is differentiable with gradient ∇f .

It is easy to see that $F(A, \cdot)$ is additively homogeneous of degree 1, so the solution p of the equation $F(A, p) = p$ may be seen as a nonlinear additive eigenvector of $F(A, \cdot)$. In this section, we give the derivative of the HOTS vector with respect to the adjacency matrix.

Proposition 7.13. *The derivative of $f \circ p$ is given by $g_{i,j} = \sum_l w_l \frac{\partial F_l}{\partial A_{i,j}}$ where*

$$w = (-\nabla f^T + (\nabla f^T e) \nabla N^T) (\nabla F - I)^\# .$$

Moreover, the matrix $(g_{i,j})_{i,j}$ has rank at most 3.

Proof. Let us differentiate with respect to $A_{i,j}$ the equation $p_l(A) = F_l(A, p(A))$. We get

$$\frac{\partial p_l}{\partial A_{i,j}} = \sum_{l \in [n]} \frac{\partial F}{\partial p_l} \frac{\partial p_l}{\partial A_{i,j}} + \frac{\partial F_l}{\partial A_{i,j}}$$

By Theorem 6.4, $\nabla F = I - \frac{1}{2} d \nabla^2 \theta$, the spectral radius is 1 and 1 is the only eigenvalue of ∇F with modulus 1.

Multiplying by $(\nabla F - I)^\#$ and using the fact that the right and left principal eigenvector of ∇F are e and the vector v such that $v_i = 1/d_{i,i}$ for all $i \in [n]$, we obtain

$$\left(I - \frac{ev^T}{v^T e}\right) \frac{\partial p}{\partial A_{i,j}} = -(\nabla F - I)^\# \frac{\partial F}{\partial A_{i,j}} \quad (7.15)$$

Multiplying by ∇N^T gives $-\frac{\nabla N^T e}{v^T e} v^T \frac{\partial p}{\partial A_{i,j}} = -\nabla N^T (\nabla F - I)^\# \frac{\partial F}{\partial A_{i,j}}$. We then use $\nabla N^T e = 1$, we reinject in (7.15) and we multiply by ∇f^T to get the result $\nabla f^T \frac{\partial p}{\partial A_{i,j}} = (-\nabla f^T + (\nabla f^T e) \nabla N^T) (\nabla F - I)^\# \frac{\partial F}{\partial A_{i,j}}$.

Finally, the equality $\frac{\partial F_l}{\partial A_{i,j}} = -\frac{1}{2} d_l \frac{\partial^2 \theta}{\partial p_l \partial A_{i,j}}$ with

$$\frac{\partial^2 \theta}{\partial p_l \partial A_{i,j}} = \frac{2\alpha - 1}{\sum_{i',j'} A_{i',j'} e^{p_{i'} - p_{j'}}} (-e^{p_i} \delta_{ij} e^{-p_j} + \delta_{li} e^{p_i} e^{-p_j} + B_l e^{p_i} e^{-p_j})$$

where $B_l = \frac{\sum_{i'} A_{i',i} e^{p_{i'} - p_l} - \sum_{j'} A_{l,j'} e^{p_l - p_{j'}}}{\sum_{i',j'} A_{i',j'} e^{p_{i'} - p_{j'}}$, shows that the matrix $(g_{i,j})_{i,j}$ has rank at most 3 since we can write it as the sum of three rank one matrices. \square

This proposition is the analog of Corollary 7.2, the latter being for Perron vector optimization problems. In both cases, the derivative has a low rank and one can compute it thanks to a Drazin inverse. Moreover, thanks to Theorem 6.4, one can apply Proposition 7.9 to ∇F . Indeed, ∇F has all its eigenvalues within $(-1, 1]$, 1 is a single eigenvalue with $ev^T / (v^T e)$ ($v_i = 1/d_{i,i}$) being the associated eigenprojector. So, for HOTS optimization problems as well as for Perron vector optimization problems, the derivative is easy to compute as soon as the second eigenvalue of ∇F is not too big.

For HOTS optimization also, we have a threshold property.

Proposition 7.14 (Threshold property). *Let A be a stationary point for the relaxed HOTS optimization problem (7.14) with associated HOTS vector p and let w be defined as in Proposition 7.13 and d as in Proposition 6.6. Let $B = \frac{\sum_{k,l} A_{k,l} e^{p_k - p_l} d_l w_l - \sum_{k,l} d_k w_k A_{k,l} e^{p_k - p_l}}{\sum_{k,l} A_{k,l} e^{p_k - p_l}}$. Then for all facultative link (i, j) , $d_j w_j > d_i w_i + B$ implies that $A_{i,j} = 1$ and $d_j w_j < d_i w_i + B$ implies that $A_{i,j} = 0$.*

Proof. The development of $\frac{\partial F_l}{\partial A_{i,j}} = -\frac{1}{2}d_l \frac{\partial^2 \theta}{\partial p_l \partial A_{i,j}}$ in the end of the proof of Proposition 7.13, shows that the derivative of the objective is given by $g_{i,j} = -\frac{1}{2} \frac{2\alpha-1}{\sum_{k,l} A_{k,l} e^{p_k - p_l}} e^{p_i - p_j} (-w_j d_j + w_i d_i + B)$.

The result follows from the fact that the problem has only box constraints. Indeed, a nonzero value of the derivative at the maximum determines whether the upper bound is saturated ($g_{i,j} > 0$) or the lower bound is saturated ($g_{i,j} < 0$). If the derivative is zero, then the weight of the link can take any value. \square

Remark 7.8. This proposition shows that $(w_i d_i)_{i \in [n]}$ gives a total order of preference in pointing to a page or another.

Proposition 7.15. *If $\tilde{\theta}$ is bounded from below, then there exists $K > 0$ such that for all n , $\|p - p_n\| \leq K \|p_{n+1} - p_n\|$.*

Proof. By Proposition 6.7, the Hessian of $\tilde{\theta}$ restricted to $H = \{y \in \mathbb{R}^n \mid \sum_{i \in [n]} y_i = 0\}$ is symmetric definite positive at any point. We thus denote

$$m = \min_{x \mid \tilde{\theta}(x) \leq \tilde{\theta}(p_0)} \lambda_{\min}(\nabla^2 \tilde{\theta}|_H)$$

m is positive because $\tilde{\theta}$ is infinitely differentiable and $\mathcal{K} = \{x \in H \mid \tilde{\theta}(x) \leq \tilde{\theta}(p_0)\}$ is compact. By Theorem 6.3, $p_n \in \mathcal{K}$ and $p_{n+1} \in \mathcal{K}$ and as \mathcal{K} is convex (because $\tilde{\theta}$ is convex), $[p_n, p_{n+1}] \in \mathcal{K}$. By comparing $\tilde{\theta}$ and the quadratic function $\phi(q) = \tilde{\theta}(p_n) + \nabla \tilde{\theta}(p_n)(p_n - q) + \frac{1}{2} \|p_n - q\|^2 m$, we deduce that $\|p_n - p\| \leq \frac{2}{m} \|\nabla \tilde{\theta}(p_n)\|$.

From the equalities obtained in Proposition 6.6 (and by a similar way as Proposition 6.6 for the one with d'), $p_{n+1} = p_n + \frac{1}{2} \log(1 - d' \nabla \tilde{\theta}(p_n)) = p_n - \frac{1}{2} \log(1 + d \nabla \tilde{\theta}(p_n))$, and $\log(1+x)^2 \geq x^2$ for $x \leq 0$, we get

$$\begin{aligned} \|p_{n+1} - p_n\|^2 &= \frac{1}{4} \sum_{i \mid \nabla_i \tilde{\theta}(p_n) \leq 0} \log(1 + d_i \nabla_i \tilde{\theta}(p_n))^2 + \frac{1}{4} \sum_{i \mid \nabla_i \tilde{\theta}(p_n) \geq 0} \log(1 - d'_i \nabla_i \tilde{\theta}(p_n))^2 \\ &\geq \frac{1}{4} \sum_{i \mid \nabla_i \tilde{\theta}(p_n) \leq 0} d_i^2 \nabla_i \tilde{\theta}(p_n)^2 + \frac{1}{4} \sum_{i \mid \nabla_i \tilde{\theta}(p_n) \geq 0} (d'_i)^2 \nabla_i \tilde{\theta}(p_n)^2 \\ &\geq \frac{1}{4} (\min_{i \in [n]} \min(d'_i, d_i)) \|\nabla_i \tilde{\theta}(p_n)\|^2 \geq (m')^2 \|\nabla \tilde{\theta}(p_n)\|^2 \end{aligned}$$

Here also $(m')^2 > 0$ because d and d' are continuous and \mathcal{K} is compact. Combining both inequalities, we get the result with $K = \frac{2}{mm'}$. \square

Corollary 7.4. *Let $(x_i)_{i \geq 0}$ be a sequence constructed by the Master Algorithm Model (Algorithm 7.3) for the resolution of the relaxed HOTS optimization problem (7.14) such that $A_n(x)$ is the Interrupted Armijo line search along the projected arc (Algorithm 7.2) and $\Delta(n) = (\Delta_0)^n$ for $\Delta_0 \in (0, 1)$. Then every accumulation point of $(x_i)_{i \geq 0}$ is a stationary point of (7.14).*

Proof. The proof follows the arguments of the proof of Theorem 7.2 but uses Proposition 7.15 instead of Proposition 7.10. \square




Figure 7.3: Strict local maximum for HOTS optimization on the small web graph of Figure 7.1. The locally optimal solution for the problem of maximizing $f(u) = \sum_{i \in I} \exp(p_i)$ presented here is to select the bold arcs with weight 1. If one replaces the arc from 20 to 17 by the arc from 17 to 20 one gets another strict local optimal solution but with a smaller value (0.166 instead of 0.169). This shows that the problem is not quasi-concave.

Example 7.2. We take the same web site as in Figures 7.1 and 7.2 with the same admissible actions. We choose the objective function $f(p) = \sum_{i \in I} \exp(p_i)$ and the normalization $N(u) = \log(\sum_{i \in I} \exp(p_i)) = 0$. The initial value of the objective is 0.142 and we present a local solution with value 0.169 on Figure 7.3.

7.8 Numerical results

By performing a crawl of our laboratory website and its surrounding pages with 1,500 pages, we obtained a fragment of the web graph. We have selected 49 pages representing a website I . We set $r_i = 1$ if $i \in I$ and $r_i = 0$ otherwise. The set of obligatory links were the initial links already present at time of the crawl, the facultative links are all other links from controlled pages except self-links.

We launched our numerical experiments on a personal computer with Intel Xeon CPU at 2.98 GHz and 8 GB RAM. We wrote the code in Matlab language.

7.8.1 HITS authority optimization

As in Section 7.5, we maximize the sum of HITS authority scores on the web site, that is we maximize $f(u) = \sum_{i \in I} r_i u_i^2$ under the normalization $N(u) = (\sum_{i \in [n]} r_i u_i^2)^{1/2} = 1$.

We use the coupled power and gradient iterations described in Section 7.4. We show the progress of the objective on Figure 7.4 and we compare coupled power and gradient iterations with classical gradient in Table 7.1.




Figure 7.4: Optimization of the sum of HITS authority scores. The figure shows that the objective is increasing during the course of the algorithm. The sum of authority values jumps from $2.1e-6$ to 0.19. Of note, the sum of authority values for a clique (all internal links activated) is $4.8e-6$. The algorithm thus gives with a small computational effort a much better solution than what is commonly considered a good strategy [LM00]. However, despite this big progress at the beginning, convergence is slow. This is a typical situation with first order methods and nonconvex optimization.

	Matrix assemblies	Power iterations	Time
Gradient (Equation (7.5))	545	-	304 s
Gradient (Remark 7.5)	324	56,239	67 s
Coupled iterations (Th. 7.2)	589	14,289	15 s

Table 7.1: Comparison of gradient algorithm with the evaluation of the gradient done by direct resolution of Equation (7.5) by Matlab “mrdivide” function, gradient algorithm with hot started power iterations described in Remark 7.5 (precision 10^{-9}) and coupled gradient and power iterations. The goal was to reach the value 0.22 on our laboratory dataset (the best value we found was 0.2285). For this problem, coupling the power and gradient iterations makes a speedup of more than four.

The best strategy of links found corresponds to a sum of HITS authority scores of 0.2285 whereas the initial situation gave a score of $3.5e-6$. This shows that HITS authority can be heavily manipulated because 49 “unknown” pages among 1,500 can get 23 % of the whole available HITS values only by changing their hyperlinks in a smart way.

The strategy has lots of binary values: only 4569 values different from 1 among 11,516 nonnegative controlled values. Moreover, the heuristic described in Section 7.5 gives a 0-1 matrix with a value at 0.07% from the weighted local optimum found. It consists in adding all possible links between controlled pages (internal links) and some external links. Following Proposition 7.11, as the controlled pages share many facultative outlinks, we can identify growing sequences in the sets of activated outlinks.

7.8.2 HOTS optimization

We consider the same website as for the numerical experiments for HITS authority. We take as objective function $f(p) = \sum_{i \in I} \exp(p_i)$ and as normalization $N(u) = \log(\sum_{i \in [n]} \exp(p_i)) = 0$.

Here, a stationary point is reached after 55 gradient iterations and 8.2 s. The initial value




Figure 7.5: Strict local optimal solution of the HOTS optimization problem. We present the adjacency matrix restricted to the set of controlled pages. There are no facultative external links added. We sorted pages by their $(w_i d_i)_{i \in [n]}$ value. The rounded black dots are the obligatory links and the blue dots are the activated facultative links. We can see that the higher the $w_i d_i$, the bigger the number of controlled pages pointing to that page and the smaller the number of pages that page points to (Proposition 7.14). It is worth noting that this local solution of the relaxed problem has only binary weights.

of the objective was 0.0198 and the final value returned by the optimization algorithm was 0.0567. The increase (186 %) is less drastic than for HITS but it is still large. The rank of the best page among the pages of the set I goes from 544th to 4th.

We give a local optimal adjacency matrix on Figure 7.5.

7.8.3 Scalability of the algorithms

We launched our algorithms on two test sets and we give the execution times on Table 7.2. Our laboratory’s website is described at the beginning of this section. The crawl of New Zealand Universities websites is available at [Pro06]. We selected the 1,696 pages containing the keyword “maori” in their url and 1,807 possible destination pages for the facultative hyperlinks, which yields 3,048,798 facultative hyperlinks. In both cases, we maximize the sum of the scores of the controlled pages.

We remark that for a problem more that 300 times greater, the computational cost does not increase that much. Indeed, the spectral gap is similar and the cost by matrix-vector product is growing only linearly thanks to the sparsity of the matrix.

7.9 Conclusion

In this chapter, we have given a scalable algorithm for the local optimization of the Perron value or a scalar function of the Perron vector of a nonnegative matrix. We have then used it to locally solve the HITS authority optimization problem and the HOTS score optimization problem.

A natural question in front of the number of existing ranking algorithms is which one is the best. First of all, the ranking should be well defined on all instances and easily computable.

	HITS (CMAP)	HOTS (CMAP)	HITS (NZU)	HOTS (NZU)
Gradient (Eq. (7.5))	1.01 s/it	0.82 s/it	out of memory	out of memory
Gradient (Rem. 7.5)	0.23 s/it	0.12 s/it	20 s/it	62 s/it
Coupled iter. (Th. 7.2)	0.038 s/it	0.04 s/it	1.0 s/it	2.9 s/it

Table 7.2: Mean execution times by gradient iterations for HITS authority and HOTS optimization for 2 fragments of the web graph: our laboratory web site and surrounding web pages (CMAP dataset, 1,500 pages) and New Zealand Universities websites (NZU dataset, 413,639 pages). The execution time with a direct resolution of Equation (7.5) by Matlab “mrdivide” function becomes very memory consuming when the size of the problem grows. It is still acceptable with matrices of size 1,500 but fails for larger matrices. The other two algorithms scale well. The execution time mainly depends on the spectral gap of the matrix, which is similar in the four cases and on the cost by matrix-vector product, which is growing only linearly thanks to the sparsity of the matrix. The number of gradient iterations is in general not dependent on the size of the problem: in our experiments for HITS authority optimization, there were even less gradient iterations with the larger dataset. When the coupled iterations is available, it gives a speedup between 3 and 20.

The classical HOTS algorithm fails for this requirement as show the experiments in Section 6.5. Then, the usual method is to define a set of web pages and to rank the pages of this test set by a jury of human experts. One compares the answer of the ranking algorithm with the man-made ranking of web pages. This test has the advantage of being compatible with any ranking algorithm and of comparing the final results. But it is an expensive test and it does not take into account the dynamical nature of the web. In particular, it does not guarantee that the chosen algorithm will be resistant to malicious behaviors.

For this, two approaches have been used. First, the design of the ranking algorithm should use a kind of mutual reinforcement property [Kle99, BP98], so that an isolated page cannot gain a large score. Second, one should test some web site patterns and see whether the answer is disproportionate. For instance, the HITS authority and hub scores are known to be exaggeratedly high on pages that are in a clique [LM00]: this is the tightly knit community effect.

The optimization procedure given in this chapter may be used as a refinement of this second approach, by testing the locally optimal web site patterns for each ranking algorithm. For instance, we have shown in Section 7.8 that the HITS authority score can still be increased a lot by adding well chosen outlinks from a clique-shaped web site. Together with the comparison of the vulnerability of the rankings to web spamming, one shall analyse the locally optimal strategies found in order to forecast what kind of strategies will be used by spammers. For instance, optimizing the PageRank (Chapter 4) or the HOTS score (Section 7.8) leads to web sites with the smallest possible number of outlinks and many inlinks while when optimizing the HITS authority score, one should also add some outlinks.

CHAPTER 8

Optimization of cancer drug treatments through age-structured population dynamics

8.1 Introduction

Up to now, we have been interested in optimization problems arising in the context of web ranking. We have thus studied the Perron vector optimization problem and a nonlinear variant of the PageRank optimization problem. In this chapter, we present and analyse a mathematical model for the optimization of cancer drug treatments in cycling cell population models with age structure. Although the application seems quite different from what we have seen in the rest of the thesis, it turns out that the models use the same fundamental tool: the growth rate of each cell population is modeled by the Perron eigenvalue of an elementwise nonnegative and sparse matrix.

Tissue proliferation in living organisms always relies on the cell division cycle: one cell becomes two after a sequence of molecular events that is physiologically controlled at each step of the cycle at so-called checkpoints [LBK⁺07, Mor06]. This process occurs in all renewing tissues, healthy or cancerous, but in cancerous tissues part of these control mechanisms are inefficient, resulting in uncontrolled tissue growth which may be given as a definition of cancer.

At the initial local stages of cancer (*i.e.*, still without more invasive processes involving tumour neoangiogenesis, digestion of the extracellular matrix and metastases), deficiencies in

the control of cell cycle checkpoints, *e.g.*, involving mutated protein p53, are the main factors responsible for this disrupted control of tissue growth.

The representation of the dynamics of the division cycle in proliferating cell by physiologically structured partial differential equations (PDEs), which dates back to McKendrick [McK26], is a natural frame to model proliferation in cell populations, healthy or cancerous. Furthermore, the inclusion in such models of targets for its physiological and pharmacological control allows one to develop mathematical methods of their analysis and control [Cla08].

Optimization of cancer therapeutics, from a clinical point of view, consists of maximizing some therapeutic index of a treatment, measuring (with various definitions) a ratio between its therapeutic outcome in terms of tumour regression and its inevitable toxic side effects on healthy tissues. It thus measures the relevance of the chosen trade-off between therapeutic efficacy and unwanted toxicity. Such an index is hence maximum when there is no toxicity and the tumour is eradicated (a most unlikely situation) and minimum when toxicity is so high that the treatment must be stopped. Giving a mathematical sense to this trade-off is the object of therapeutic optimization by mathematical methods [AGLG11, ALG07b, ALG09, BCL06, Cla07, Cla09, Cla11].

Optimizing anticancer treatments may plainly consist of defining the drug delivery time schedule (or, rather, its output on p53 or CDK inhibitor concentration at the target cell population level) that will kill as many cancer cells as possible, and then adapt it by trials and errors to clinical tolerability constraints. But since we know that drugs act on proliferating healthy cells (in fast renewing tissues, such as gut epithelium or bone marrow) by the same mechanisms as in cancer drugs, and that these unwanted side-effects on healthy cells are actual clinical issues that limit the use of these drugs, it is legitimate to consider the drug delivery optimization problem as a problem of *optimization under constraints*, the solution of which is a trade-off between the objective function (decreasing the cancer cell population) and the constraint (maintaining the healthy cell population over a tolerable level).

Placing ourselves within the modelling frame of a generic drug acting in parallel on a healthy tissue and on a tumour, both represented by cell population dynamic systems of equations, it has already been shown that it was possible, using a simplified set of ordinary differential equations (ODEs) with a physiological circadian clock control on the pharmacodynamics of a one-drug external control, to obtain optimized drug delivery schedules, solutions to an optimization problem under constraints [BCL06, Cla07]. These schedules are optimal in the sense that, constraining the healthy cell population to remain over an absolute tolerability threshold (absolute, but adaptable to what could be the patient's state of health), the tumour cell kill is maximized by a delivery drug flow that takes into account the effects of circadian clocks on both cell populations.

The assumption used there was that there was a best time to kill tumour cells, which was at the same time the best to preserve healthy cells from toxicity. Although it has found some experimental support explored in [BCL06, Cla07], this assumption may be forsaken to comply with the more likely conjecture of a clear obedience of healthy proliferating cells to circadian clock synchronizing messages, and a looser obedience to, or total ignorance of, the same messages in tumour cells. It is this assumption that we will now put forward, as it has also been used by others in different modelling settings [AGLG11, ALG07b, ALG09].

In this chapter, we present and analyse a mathematical model for the optimization of cancer drug treatments in cycling cell population models with age structure. We consider a drug, 5-FluoroUracil (5-FU), that prevents cells from starting mitosis.

The proliferating healthy and cancer cell populations are represented by the same age-

structured model of the McKendrick type, with different physiological controls for the two cases. Each dynamics is given by physiologically structured PDEs where the dynamic variables are the number of cells of each age in each phase of the proliferation cycle: first growth phase (G_1), DNA synthesis followed by second growth phase (S/G_2) and mitosis (M), and the only structure variable is age in the phase. The parameters of the model are the death rates and transitions from a phase to the next in the cell cycle.

In this work, we assume that proliferation is physiologically controlled by a circadian clock [L02, L06, L08, LACG08, LOD⁺10, LS07b], which implies that the coefficients of the model are not only age and phase-dependent but also 24h-periodic functions. Our fundamental hypothesis is that healthy and cancer cells proliferate following the same model, but that cancer cells are characterized by a looser response to the circadian control, which gives them a proliferation advantage [AGLG11, ALG07b, ALG09].

We show how recent fluorescence-based image modelling techniques performed at the single cell level in proliferating cell populations allow one to identify some key parameters of the population dynamics, giving the transition rates from a phase to the next one when there is no circadian control. Then, we inferred the transition rates with circadian control by assuming that the actual transition rate is the product of the rate without circadian control and of a 24h-periodic physiological control.

Then, we consider time-dependent 5-FU infusions that disrupt the transition from phase G_2 to M . We study the problem of minimizing the growth rate of the cancer cell population, modeled by the Floquet eigenvalue of the population dynamics, with the constraint that the growth rate of the healthy cell population remains over a given toxicity threshold. The goal is to find (periodic) chemotherapy schedules that are viable in the long term and effective in the fight against cancer.

When we discretize the problem, the Floquet eigenvalues are approached by the Perron eigenvalues of sparse nonnegative matrices. We developed a multiplier's method for the local optimization of the growth rates, that takes advantage of a low rank property of the gradient of the Perron eigenvalue. The eigenvalue optimization algorithm is based on the algorithm developed in Chapter 7. We calculated the gradient of the objective function at stake and adapted the coupling of power and gradient iterations to the multiplier's method. This chapter relies on the work presented in [BCF⁺11a, BCF⁺11b, BCF12].

8.2 Drugs used in cancer treatments and their targets

8.2.1 Fate of drugs in the organism: molecular pharmacokinetics-pharmacodynamics

Anticancer drugs are delivered into the general circulation, either directly by intravenous infusion, or indirectly by oral route, intestinal absorption and enterohepatic circulation (*i.e.*, entry in the general blood circulation from the intestine via the portal vein towards the liver, and possibly back from the liver to the intestine via bile ducts). Their fate, from introduction in the circulation until presence of an active metabolite in the neighborhood of their intracellular targets, can be represented by *pharmacokinetic (PK)* compartmental ordinary differential equations (ODEs) for their concentrations. It is also theoretically possible to represent this fate by spatial partial differential equations (PDEs) with boundary conditions instead of exchange rules between compartments when data on spatial diffusion of the drugs

and some geometry of their distribution domain is known - but this is seldom the case.

Then, in the cell medium, either an individual cell, or a mean intracellular medium in a population of cells, *pharmacodynamic (PD)* differential equations must be used to relate local drug concentrations with molecular effects on their targets. At this level of description, it is *a priori* more relevant to describe by physiologically structured than by spatially structured models the population of cells under pharmacological attack, since anticancer drugs act mainly by blocking the cell division cycle, which does not give rise to a spatially structured cell population (apart from the very early stages of avascular spheroid tumour growth, little geometry is relevant to describe a tumour seen under the microscope).

8.2.2 Cytotoxics and cytostatics

Driving cells to their death may be obtained either by damaging the genome, or more indirectly by impairing essential mechanisms of the cell division cycle, such as enzymes thymidylate synthase (an enzyme that plays an essential role in DNA synthesis and is one of the main targets of cytotoxic drug 5-FU) or topoisomerase I (another essential enzyme of DNA synthesis, target of cytotoxic drug irinotecan). The resulting damaged cell, unable to proceed until division into two viable cells, is normally blocked at one or the other checkpoint, mainly G_1/S or G_2/M (recall that the cell division cycle is classically divided into 4 successive phases, G_1 , S for DNA synthesis, G_2 , and M for mitosis).

Then, unless it may be repaired by specific enzymes - that are often overexpressed in cancer cells -, these impaired cells, blocked at a checkpoint, are subsequently sent to ‘clean death’ by the physiological mechanism of apoptosis (also possibly impaired in cancer, resulting in abnormal cells bypassing these checkpoints). As mentioned above, we define here this class of drugs, that have for their ultimate mission to kill cancer cells - even if their primary action is not to directly damage the genome, but rather to damage cell cycle enzymes - as *cytotoxics*.

For our numerical experiments, we chose to consider 5-FluoroUracil (5FU) infusions. This anticancer drug has an S phase specificity, acting via DNA damage and involving a preserved p53 protein control (having nevertheless in mind that p53 is mutated in many cancers [KB04, VLL00]). When 5-FU is infused in the body of the patient, some cells will have damaged DNA and will be blocked at the checkpoint between phases G_2 and M . We assume that DNA repair mechanisms exist and thus that when the drug infusions stop, the cells can go on proliferating.

We reserve the term *cytostatic* to those non cell-killing drugs that merely slow down proliferation, usually by maintaining cells in G_1 with possible exit to G_0 , that is by definition the quiescent phase, i.e., the subpopulation of cells that are not committed in the cell division cycle. Indeed, before the restriction point inside G_1 , cells may stop their progression in the cell cycle and go back to quiescence in G_0 . This last category comprises all drugs that act as antagonists of growth factors, which may be monoclonal antibodies or tyrosine kinase inhibitors.

8.2.3 Representation of drug targets

It is appropriate to consider anticancer drugs, cytotoxic or cytostatic, and their targets, through their effects on the cell cycle in cell populations. This of course assumes that a model of the cell cycle in a proliferating cell population is given.

In a review article [KS06], Kimmel and Świerniak considered two possibilities to represent in a mathematical model the action of cytotoxic drugs on their targets in a proliferating cell population: either by a possible direct effect on cell death, enhancing it, presumably by launching or accelerating the apoptotic cascade in one or more phases of the cell cycle, or by a blockade of one or more transitions between two phases, arresting the cycle at some checkpoint, most often with the involvement of protein p53, and only secondarily launching cell death.

This is indeed a general alternative in the representation of the effects of cytotoxics. If no cell cycle phase structure has been put in the population dynamic model used to represent the evolution of the cell populations at stake, i.e., when no account is taken of cell cycle phases in these populations, then only the first possibility exists: modulation of a death term.

As regards cytostatics (which by definition are not supposed to kill cells, at least not directly), the representation of their action in physiological models with age structure for the cell cycle should be somewhat different. It can be done either by a slowdown of the progression speed in the G_1 phase (or in the proliferating phase in a one-phase model) or by an action on the exchanges between non proliferating (G_0) and proliferating phases when a G_0 phase is represented in the model.

It is also possible to combine cytostatic and cytotoxic effects in the same model. In [HWAW07], for instance, the authors use an age-structured model with a 1-phase proliferative subpopulation exchanging cells with a nonproliferative cell compartment to combine a slowdown effect on proliferation for the cytostatic effect with an increase in the cell death term for the cytotoxic effect - of the same drug, lapatinib, a tyrosine kinase inhibitor, in their case, the variation between these effects depending on the dose. Acting on two different targets in a cell cycle model by two different drugs, a cytotoxic and a cytostatic one, in the same cell population is thus possible, and such models are thus amenable to study and optimise combination therapies, such as cetuximab+irinotecan advocated in [CHS⁺04].

8.3 Age-structured linear models for tissue proliferation and its control

8.3.1 Linear and nonlinear models

Physiologically structured cell population dynamics models have been extensively studied in the last 25 years, see *e.g.* [Ari95, AK93, AS97, MD86, BCF12]. We consider here a typically age-structured cell cycle model, since our aim was to represent the action of cytotoxic anticancer drugs, which always act onto the cell division cycle in a proliferating cell population. The model chosen, of the McKendrick type [McK26], is linear. This may be considered as a harsh simplification to describe biological reality, which involves nonlinear feedbacks to represent actual growth conditions such as population size limitation due to space scarcity. Nonetheless, having in mind that linear models in biology are just linearizations of more complex models (for instance considering the fact a first course of chemotherapy will most often kill enough cells to make room for a non space-limited cell population to thrive in the beginning) we think that it is worth studying population growth and its asymptotic behavior in linear conditions and thus analyse it using its growth (or Malthus) exponent. This first eigenvalue of the linear system may be considered as governing the asymptotic behavior, at each point where it has been linearized, of a more complex nonlinear system, as described in

[BBCP08, BBCRB08].

8.3.2 Age-structured models for tissue proliferation and its control

We know that circadian clocks ([L02, L06, L08, LACG08, LOD⁺10, LS07b], Section 1.2) normally control cell proliferation, by gating at checkpoints between cell cycle phases (i.e., by letting cells pass to the next phase only conditionally). We also know that circadian clock disruption has been reported to be a possible cause of lack of physiological control on tissue proliferation in cancer [LOD⁺10], a fact that we will represent in our model to distinguish between cancer and healthy cell populations.

The representation of the dynamics of the division cycle in proliferating cell proliferations by physiologically structured partial differential equations (PDEs) is thus a natural frame to model proliferation in cell populations, healthy or cancerous. The inclusion in such proliferation models of targets for its control, physiological (circadian) and pharmacological (by drugs supposed to act directly on checkpoints), allows to develop mathematical methods of their analysis and therapeutic control [BCF⁺11a, BCF⁺11b, Cla08], in particular for cancer chronotherapeutics, i.e., when the drug control is made 24h-periodic to take advantage of favorable circadian times.

We consider here an age-structured cell cycle model, in which the cell division cycle is divided into I phases (classically 4: G_1, S, G_2 and M), and the variables are the densities $n_i(t, x)$ of cells having age x at time t in phase i . Equations read

$$\begin{cases} \frac{\partial n_i(t, x)}{\partial t} + \frac{\partial n_i(t, x)}{\partial x} + d_i(t, x)n_i(t, x) + K_{i \rightarrow i+1}(t, x)n_i(t, x) = 0, \\ n_{i+1}(t, 0) = \int_0^\infty K_{i \rightarrow i+1}(t, x)n_i(t, x)dx, \\ n_1(t, 0) = 2 \int_0^\infty K_{I \rightarrow 1}(t, x)n_I(t, x)dx. \end{cases} \quad (8.1)$$

Together with an initial condition $(n_i(t = 0, \cdot))_{1 \leq i \leq I}$. This model was first introduced in [CLMP03]. The case $I = 1$ has received particular attention [CGL09, CGL11]. In this model, in each phase, the cells are ageing with constant speed 1 (transport term), they may die (with rate d_i) or go to next phase (with rate $K_{i \rightarrow i+1}$) in which they start with age 0. We write it in its highest generality. If we want to represent the effect of circadian rhythms, we usually consider time-periodic coefficients, the period being of course $24h$.

Note that we will not consider here a resting phase (G_0), with exchanges with the G_1 phase. It is certainly possible to do it in a biologically more realistic setting, since even in fast renewing tissues, all cells are not in a proliferative state [BT70, VVB03], and also since it has been done already [BBCP08, BBCRB08, BPMM03], even in a completely linear setting [GGTW11] (i.e., without nonlinear feedback). Nevertheless, since the observations on the basis of which we performed the identification of our model parameters (see below Section 8.5.1) report recordings on proliferating cells only, we have limited our theoretical frame here to the classical G_1, S, G_2 and M phases.

8.3.3 Basic facts about age-structured linear models

One of the most important facts about linear models is its trend to exponential growth. Solutions to (8.1) satisfy (if the coefficients are time-periodic, or stationary) $n_i(t, x) \sim$

$C^0 N_i(t, x)e^{\lambda t}$, where N_i are defined by (for T -periodic coefficients)

$$\begin{cases} \frac{\partial N_i(t, x)}{\partial t} + \frac{\partial N_i(t, x)}{\partial x} + (\lambda + d_i(t, x) + K_{i \rightarrow i+1}(t, x))N_i(t, x) = 0, \\ N_{i+1}(t, 0) = \int_0^\infty K_{i \rightarrow i+1}(t, x)N_i(t, x)dx, \\ N_1(t, 0) = 2 \int_0^\infty K_{I \rightarrow 1}(t, x)N_I(t, x)dx, \\ N_i > 0, \quad N_i(t + T, \cdot) = N_i(t, \cdot), \quad \sum_i \int_0^T \int_0^\infty N_i(t, x)dxdt = 1. \end{cases} \quad (8.2)$$

The study of the growth exponent λ , first eigenvalue of the system, thus governs the long-time behavior of the population (since the N_i are bounded) and is therefore of crucial importance. For further reading about the asymptotic behavior of (8.1), the reader may consult [Per07] (chapter 3) for an overview of the subject.

In the 1-phase case ($I = 1$), *i.e.*, if only the total cell cycle duration is taken into consideration, extended studies of the first and second eigenvalues (there is only one positive eigenvalues, but others exist, that are complex) of the system have been performed in [CMTU01], following [MD86], with proposed experimental methods using flow cytometry to identify these eigenvalues. In the present study, we will not consider other eigenvalues than the first one, but we are fully aware of the fact that considering the second eigenvalue (its real and imaginary parts) may be of importance if one wants to precisely describe in particular transient phenomena that appear when control changes occur at cell cycle phase transitions.

We focus now on the case of stationary phase transition coefficients ($K_{i \rightarrow i+1}(t, x) = K_{i \rightarrow i+1}(x)$) and we do not consider death rates ($d_i = 0$). Note that if one considers constant nonzero death rates, the problem does not change, only the eigenvalue λ is then in fact $\lambda + d$, as one can see in the equations of system (8.2). As shown in [CLMP03], the first eigenvalue λ is then solution of the following equation, which in population dynamics is referred to, in the 1-phase case ($I = 1$) with no death term, as Euler-Lotka's equation

$$\frac{1}{2} = \prod_{i=1}^I \int_0^{+\infty} K_{i \rightarrow i+1}(x) e^{-\int_0^x K_{i \rightarrow i+1}(\xi) d\xi} e^{-\lambda x} dx. \quad (8.3)$$

Integrating the first equation of System (8.1) along its characteristics [Per07], we can in the stationary case with no death rate derive the formula

$$n_i(t + x, x) = n_i(t, 0) e^{-\int_0^x K_{i \rightarrow i+1}(\xi) d\xi}.$$

This can be interpreted in the following way: the probability for a cell which entered phase i at time t to stay for at least an age duration x in phase i is given by

$$P(\tau_i \geq x) = e^{-\int_0^x K_{i \rightarrow i+1}(\xi) d\xi}.$$

The time τ_i spent in phase i is thus a random variable on \mathbb{R}_+ , with probability density function f_i given by

$$dP_{\tau_i}(x) = f_i(x)dx = K_{i \rightarrow i+1}(x) \cdot e^{-\int_0^x K_{i \rightarrow i+1}(\xi) d\xi} dx,$$

or equivalently:

$$K_{i \rightarrow i+1}(x) = \frac{f_i(x)}{1 - \int_0^x f_i(\xi) d\xi} \quad (8.4)$$

8.3.4 Discretization scheme

In this section, we come back to the general case where $K_{i \rightarrow i+1}$ depends on age and time and the death rate is nonzero. In order to simulate the evolution of the cancer and healthy cell populations, we propose a discretization scheme. For numerical convenience, we had to suppose that the cells can not grow beyond a maximal age. Indeed, numerically assessing this assumption, we came to the conclusion that it had little influence on the value of $\lambda > 0$ provided that this maximal age was large enough, so that, we arbitrarily took a maximal age of 10 days for each one of the phases. This age limit can be interpreted as the time from which DNA damages are no longer repaired and the cells are sent to apoptosis.

We base our discretization scheme on the one presented in [CGL09]. We denote by Δt and Δx respectively the time and age steps, *i.e.* $t = k\Delta t$ and $x = j\Delta x$ where $k = 0 \dots \lfloor \frac{T_{tot}}{\Delta t} \rfloor$ and $j = 0 \dots \lfloor \frac{X_{tot}}{\Delta x} \rfloor$, $\lfloor z \rfloor$ representing the integer part of the real number z , and T_{tot} and X_{tot} being respectively the total time of the simulation and the maximal time a cell can spend in one phase of the cell cycle. For $k \in \mathbb{N}$, $j \in \mathbb{N}$ and for $i = 1, \dots, I$, we consider the following quantities:

$$\begin{aligned} n_i^{k,j} &= n_i(k\Delta t, j\Delta x) \\ K_{i \rightarrow i+1}^{k,j} &= K_{i \rightarrow i+1}(k\Delta t, j\Delta x) \\ d_i^{k,j} &= d_i(k\Delta t, j\Delta x) \end{aligned}$$

We use a first order finite difference scheme on the I -phase mathematical model. Assuming $\Delta t = \Delta x$ (CFL=1), we have:

$$\begin{aligned} n_i^{k+1,j} &= \frac{n_i^{k,j-1}}{1 + \Delta t(K_{i \rightarrow i+1}^{k+1,j} + d_i^{k+1,j})} & i = 1 \dots I, \quad k = 0 \dots k_{max}-1, \quad j = 1 \dots j_{max} \\ n_i^{0,j} &= n_i^0(j\Delta x) & i = 1 \dots I, \quad j = 0 \dots j_{max} \\ n_1^{k+1,0} &= 2 \sum_{j=0}^{j_{max}} \frac{\Delta t K_{I \rightarrow 1}^{k+1,j+1}}{1 + \Delta t(K_{I \rightarrow 1}^{k+1,j+1} + d_I^{k+1,j+1})} n_I^{k,j} & k = 0 \dots k_{max}-1 \\ n_i^{k+1,0} &= \sum_{j=0}^{j_{max}} \frac{\Delta t K_{i-1 \rightarrow i}^{k+1,j+1}}{1 + \Delta t(K_{i-1 \rightarrow i}^{k+1,j+1} + d_{i-1}^{k+1,j+1})} n_{i-1}^{k,j} & i = 2 \dots I, \quad k = 0 \dots k_{max}-1 \end{aligned}$$

where $k_{max} = \lfloor \frac{T_{tot}}{\Delta t} \rfloor$, $j_{max} = \lfloor \frac{X_{tot}}{\Delta x} \rfloor$, n_i^0 denoted the initial density of cells in each of the two phases.

Unlike the scheme presented in [CGL09], this scheme ensures mass conservation of cells through the phase transitions for any discretization step. Indeed, for $i \leq I-1$, we have the relation

$$\sum_{j \geq 0} n_i^{k,j} = n_{i+1}^{k+1,0} + \sum_{j > 0} n_i^{k+1,j} + \sum_{j \geq 0} \frac{\Delta t d_i^{k+1,j+1}}{1 + \Delta t(K_{i \rightarrow i+1}^{k+1,j+1} + d_i^{k+1,j+1})} n_i^{k,j} + \frac{n_i^{k,j_{max}}}{1 + \Delta t(K_{i \rightarrow i+1}^{k+1,j_{max}+1} + d_i^{k+1,j_{max}+1})},$$

which means that cells at time k either go to next phase or get older or die or go beyond the age limit and are thus “killed.” This property gives a better coherence with biological reality.

8.3.5 Numerical determination of the growth exponent

If transition rates do not depend on time, which is the case when there is no circadian control, the first eigenvalue λ of the system (8.2) is given as the only positive solution to Equation (8.3). In the general case, transition rates depend on time. Insofar as we are interested in the effect of circadian rhythms on cell cycle progression, we restrict this study to time-periodic transition rates, the period being equal to $T = 24$ hours. Over a time period, let $n^k = (n_1^{k,0}, n_1^{k,1}, \dots, n_1^{k,j_{max}}, n_2^{k,0}, \dots, n_2^{k,j_{max}})$ be a vector in $\mathbb{R}^{2(j_{max}+1)}$. As in the case of the 1-phase model presented in [CGL09], we can write $n^{k+1} = M_k n^k$ where

$$M_k = \begin{pmatrix} A1_k & 0 & \dots & 0 & BI_k \\ B1_k & A2_k & 0 & \dots & 0 \\ 0 & B2_k & A3_k & 0 & 0 \\ \vdots & \dots & \ddots & \ddots & 0 \\ 0 & \dots & \dots & B(I-1)_k & AI_k \end{pmatrix}$$

The blocks Ai_k and Bi_k are $(j_{max} + 1) \times (j_{max} + 1)$ matrices defined by:

$$Ai_k = \begin{pmatrix} 0 & \dots & \dots & 0 & 0 \\ \frac{1}{1+\Delta t(K_{i \rightarrow i+1}^{k+1,1} + d_i^{k+1,1})} & & & \vdots & \vdots \\ 0 & \ddots & & \vdots & \vdots \\ \vdots & & \ddots & 0 & 0 \\ 0 & \dots & 0 & \frac{1}{1+\Delta t(K_{i \rightarrow i+1}^{k+1,j_{max}} + d_i^{k+1,j_{max}})} & 0 \end{pmatrix} \quad i = 1 \dots I$$

$$Bi_k = \begin{pmatrix} \frac{\Delta t K_{i \rightarrow i+1}^{k+1,1}}{1+\Delta t(K_{i \rightarrow i+1}^{k+1,1} + d_i^{k+1,1})} & \dots & \frac{\Delta t K_{i \rightarrow i+1}^{k+1,j_{max}+1}}{1+\Delta t(K_{i \rightarrow i+1}^{k+1,j_{max}+1} + d_i^{k+1,j_{max}+1})} \\ 0 & \dots & 0 \\ \vdots & & \vdots \\ 0 & \dots & 0 \end{pmatrix},$$

$$\text{and } BI_k = \begin{pmatrix} \frac{2\Delta t K_{I \rightarrow 1}^{k+1,1}}{1+\Delta t(K_{I \rightarrow 1}^{k+1,1} + d_I^{k+1,1})} & \dots & \frac{2\Delta t K_{I \rightarrow 1}^{k+1,j_{max}+1}}{1+\Delta t(K_{I \rightarrow 1}^{k+1,j_{max}+1} + d_I^{k+1,j_{max}+1})} \\ 0 & \dots & 0 \\ \vdots & & \vdots \\ 0 & \dots & 0 \end{pmatrix}.$$

This matrix M_k depends only on k and is T -periodic. Moreover, if we define the matrix $\mathbb{M} = M_{N_T} M_{N_T-1} \dots M_2 M_1$ where $N_T = T/\Delta t$, we have $n^{N_T} = \mathbb{M} n^0$. The following proposition is simple adaptation of its analogue in [CGL09].

Proposition 8.1. *If for all phases i and for all $t \in [0; T_{tot}]$, there exists $x_0 \in [0; X_{tot}]$ such that $K_{i \rightarrow i+1}(x, t) > 0$ for all $x \geq x_0$, then $\mathbb{M} = M_{N_T} M_{N_T-1} \dots M_2 M_1$ is a nonnegative and irreducible matrix.*

By the Perron-Frobenius theorem (see [BP94] for instance), this proposition means that \mathbb{M} has a simple and positive principal eigenvalue ρ associated with nonnegative left and right eigenvectors unique up to a given normalization. We calculated this eigenvalue ρ by means of the power algorithm [GL89]. The Floquet eigenvalue, or growth exponent, λ can then be approximated by $\frac{1}{T} \log(\rho)$ [CGL09].

8.4 The control and its missions: representing the action of drugs

In the previous section, we have described a dynamic model for cell populations. This model can thus be seen as a controlled dynamic system with drug effects as their control functions. Various examples of such drug effects have been given in Section 8.2. Introducing pharmacokinetics (i.e., evolution of concentrations) for the drugs chosen produces additional equations to the cell population dynamic model, and their pharmacodynamics (i.e., actual drug actions) modify this cell dynamics according to the target and to the effect of the drugs. Then, optimization of cancer treatments can be represented as an optimal control problem on this controlled dynamic system. In this section, we first discuss how the drug infusions are taken into account in the model, then we give examples of objective functions and constraints considered in the literature on the treatment of cancers.

8.4.1 The control functions

In the population dynamics presented in Section 8.3, the dynamic variables are the number of cells of each age in each phase of the proliferation cycle: first growth phase (G_1), DNA synthesis followed by second growth phase (S/G_2) and mitosis (M). They are functions depending on the age in the phase (the structure variable). We chose to merge phases S and G_2 because the identification procedure that we will present in Section 8.5 only differentiates the phase G_1 from the others. We then assumed that the mitosis had a fixed duration.

The control represents the action of the 5-FU infusions on the death rates and transitions from a phase to the next in the cell cycle. We assume that these infusion have no direct impact on the death rate but that they disrupt the natural transition rate from phase S/G_2 to phase M , $K_{2 \rightarrow 3}^0(t, x)$ in the following way

$$K_{2 \rightarrow 3}(t, x) = K_{2 \rightarrow 3}^0(t, x)(1 - g_2(t)) . \quad (8.5)$$

$g_2(t)$ is the effect at the cell level of the drug infusion at time t on the transition rate from phase S/G_2 to phase M . No drug corresponds to $g_2(t) = 0$ ($g_1(t) = g_3(t) = 0$ since the other transitions are not affected by 5-FU infusions), a transition-blocking infusion corresponds to $g_2(t) = 1$. As $K_{2 \rightarrow 3}^0(t, x)$ is 24h-periodic, if $g_2(t)$ is also 24h-periodic, then $K_{2 \rightarrow 3}(t, x)$ is also 24h-periodic and we can define the Floquet eigenvalue of the dynamics following (8.2). We thus describe long-range drug treatment in a stationary state by controlling growth exponents, i.e., first eigenvalues of both cell population systems, healthy and cancer, simultaneously. It may be possible to consider drug infusion schedules with longer periods but numerical experiments suggest that optimal drug infusion patterns are 1-day periodic, so we restrict to 1-day periodic drug infusion strategies.

Alternatively, instead of a control function, one may consider simpler predefined infusion schemes with only a small number of control parameters. Such infusion schemes may represent either a simple model for an early study or a consequence of technical constraints such as the fact that oral drugs can only be administered at fixed hours (at meal time for instance). Examples of such parameters are the period of a periodic scheme [Pan97, Web90] or the phase difference between a circadian clock and the time of drug infusion initiation [ALG09].

8.4.2 Objective functions: measuring the output

An optimization problem consists in maximizing or minimizing a given real-valued objective function, that models the objective we want to reach.

The main purpose of a cancer treatment is to minimise the number of cancer cells. When the model takes into account the number of cancer cells directly [ALG09, BCL06, dLMS09, FP00, KS06, LMS11, PPdS95, VOA10], the objective function is simply the value of the coordinate of the state variable corresponding to the number of cancer cells at a time T , T being either fixed or controlled.

The optimization problem can also be formalized as the minimization of the asymptotic growth rate of the cancer cell population [PA95, Web90]. Hence, the number of cancer cells will increase more slowly, or even eventually decrease. We will present this approach in a linear frame (hence controlling eigenvalues) in Section 8.4.4.

8.4.3 Constraints, technological and biological, static or dynamic

8.4.3.a Toxicity constraints A critical issue in cancer treatment is due to the fact that drugs usually exert their effects not only on cancer cells but also on healthy cells. A simple way to minimise the number of cancer cells is to deliver a huge quantity of drug to the patient, who is however then certainly exposed at high lethal risk. In order to avoid such “toxic solutions”, one may set constraints in the optimization problem, which thus becomes an optimization problem under constraints.

Putting an upper bound on the drug instantaneous flow [FP00] and/or on the total drug dose is a simple way to prevent too high a toxicity for a given treatment. A bound on total dose may also represent a budget limit for expensive drugs [LMS11].

However, fixed bounds on drug doses are not dynamic, i.e., they do not take into account specificities of the patient’s metabolism and response to the treatment, other than by adapting daily doses to fixed coarse parameters such as body surface or weight (as is most often the case in the clinic so far). In order to get closer to actual toxicity limits, and hoping for a better result, it is possible to consider instead a lower bound on the number of healthy cells, as in [BCL06]. In the same way, using a Malthusian growth model, where growth exponents are the targets of control, such a constraint becomes a lower bound on the asymptotic growth rate of the healthy cell population [Web90].

In the same way, a drug used in a treatment must reach a minimal concentration at the level of its target (which blood levels reflect only very indirectly) to produce therapeutic effects. Classically, clinical pharmacologists are accustomed to appreciating such efficacy levels by lower threshold blood levels, that are themselves estimated as functions of pharmacokinetic parameters such as first and second half-life times and distribution volume of the drug, with confidence interval estimates for a general population of patients. As in the case of toxicity, a more dynamic view is possible, by considering drug levels that decrease the number of cancer cells, that is, which yield a negative growth rate in the cancer cell population.

This leads to the definition of admissible sets for drug infusion flows, the union of $\{0\}$ and of a therapeutic range containing the infusion levels that are at the same time efficient and not too toxic (such a constraint is considered in [VOA10]). Those admissible sets are rather difficult to take into account, however, as they lead to complex combinatorial problems.

An approach that is consequently often chosen (see [LMS11] for instance) is to forget this constraint in the model and to *a posteriori* check that the optimal drug infusion schedules found are high enough to be efficient when they are nonzero.

That may be an elementary reason why so called bang-bang controls (i.e., all-or-none) are of major interest in chemotherapy optimization: they are defined as controls such that at each time, either the drug infusion flow is the smallest possible (i.e., 0), or it is the highest

possible. Even though it is now easy to use in the clinic (and also in ambulatory conditions) programmable pumps that may deliver drug flows according to any predefined schedule with long-lasting autonomy, solutions to optimization problems often turn out to be bang-bang (tap open-tap closed).

But solutions to optimization problems in cancer chemotherapy are not always bang-bang, when considerations other than on simple parallel growth of the two populations are taken into account, and this includes competition, when the two populations are in contact, e.g. in the bone marrow normal haematopoietic and leukaemic cells, or when both populations are submitted to a common - but differently exerted - physiological control, such as by circadian clocks [BCL06].

Another interesting approach, relying on two models, one of them including the cell division cycle [PA95], and putting the optimal control problem with toxicity constraints, is developed in [DDP08]. The optimal control problem is solved by using the industrial software gPROMS[®].

8.4.3.b Drug resistance Whereas therapeutic efficacy and limitation of toxic adverse effects are the first concern when dealing with chemotherapy, the frequent development of drug resistances in the target cancer cell populations is certainly the second bigger issue in the clinic. The development of such resistances may come from overexpression in individual cells of defense mechanisms as an exaggeration of physiological phenomena, such as are ABC transporters (the P-gp, or P-glycoprotein, being its most known representant), but they may also result, at least as likely, in proliferative populations encompassing mitoses, from mutations yielding more fit, i.e., resistant in the presence of drug, subpopulations.

A classical solution to this problem is to forbid too low drug concentrations, that are supposed to create environmental conditions favorable to the development of more fit drug resistant cell populations without killing them, as is also the case, for instance, in antibiotherapy with bacteria. Nevertheless, other, more recent, arguments to support an opposite view, have been put forth: assuming that there exists a resistant cell population at the beginning of the treatment, or that it may emerge during the treatment, then delivering high drug doses often produces the effect to kill all sensitive cells, giving a comparative fitness advantage to resistant cells, that subsequently become very hard to eradicate. Thus a paradoxical solution has been proposed, at least in slowly developing cancers: killing just enough cancer cells to limit tumour growth, but letting enough of these drug sensitive cancer cells to oppose by competition for space the thriving of resistant cells, that are supposed to be less fit, but just the same, usually slowly, will invade all the tumour territory if no opponents are present [Gat09, GSGF02]. Indeed, such free space left for resistant tumour cells to thrive, when high drug doses have been administered with the naive hope to eradicate all cancer cells, may result in the rise of tumours that escape all known therapeutics, a nightmare for physicians which is unfortunately too often a clinical reality. Hence the proposed strategy to avoid high doses, that are able to kill all sensitive cells, and to only *contain* tumour growth by keeping alive a minimal population of drug-sensitive tumour cells.

Both those constraints, toxicity and resistance, can be considered as part of the objective function by setting the objective to be a balance between two objectives. For instance, Kimmel and Swierniak in [KS06] proposed to minimise a linear combination of the number of cancer cells and of the total drug dose. This yields an unconstrained optimization problem, that has a simpler resolution, while still taking into account the diverging goals of minimizing the

number of cancer cells and keeping the number of healthy cells high enough.

But whereas cancer and healthy cells are two quite distinct populations, with growth models that may easily be distinguished and experimentally identified by their parameters, it is more difficult to take into account the evolutionary lability (i.e., the genomic instability) and heterogeneity of cancer cell populations with respect to mutation-selection towards drug resistance, according to evolution mechanisms that are not completely elicited. Note that *acquired* (as opposed to *intrinsic*, i.e., genetically constitutive) drug resistance may result as well from individual cell adaptation (enhancement of physiological mechanisms) as from genetic mutations, both under the pressure of a drug-enriched environment, as discussed in [CLP11]. In this respect, acquired resistance may be reversible, if no mutation has initiated the mechanism, or irreversible, and it is likely irreversible in the case of intrinsic resistance.

Ideally, the optimal solution of a therapeutic control problem should take into account both the drug resistance (using evolutionary cell population dynamics) and the toxicity constraints, but these constraints have usually been treated separately so far. Whereas the difficult problem of drug resistance control is certainly one of our concerns in a cell Darwinian perspective, in the sequel we shall present only results for the (easier) toxicity control problem.

8.4.4 An eigenvalue optimization problem with constraints

In a cancer chemotherapy focus, we propose to minimize the growth rate of the cancer cell population while maintaining the growth rate of the healthy cell population above a threshold Λ . Infusion here may be thought of as referring to the drug 5-FluoroUracil (5FU) that acts via DNA damage, thus directly blocking G_2/M transition, as explained in (8.5).

We consider two cell populations with their respective dynamics. We modeled both of them with an age-structured cell population model (8.1) but with different parameters. In fact, we assume that the transition and death rates are the same when there is no circadian control but that cancer cells have a looser answer to the physiological circadian control. We will give the precise value of these parameters in Section 8.5 below. We assumed that the drug has the same effect on both populations, which couples their behaviors through the drug infusions.

Then we model the growth rate of both populations by the Floquet eigenvalue of the corresponding model (λ_C for cancer cells and λ_H for healthy cells). We obtain the following Floquet eigenvalue optimization problem with constraints:

$$\begin{aligned} \min_{g(\cdot)} \lambda_C(g) \\ \lambda_H(g) \geq \Lambda \\ g \text{ 24h-periodic} \end{aligned} \tag{8.6}$$

8.5 Identification of model parameters: the target model and drug effects

For the pharmacokinetic and pharmacodynamic models, we simply assume that the drug infusions immediately alter the transition rates with an affine dependency such that with the maximal infusion level, the transition rate is minimal. We nevertheless propose a precise

identification of parameters of the population dynamics model, the most important of which are the transition rates. The identification procedure relies on Eq. (8.4), that gives a relation between the transition rate $K_{i \rightarrow i+1}$ and the probability law of the duration of phase i .

8.5.1 FUCCI

From a biological point of view, the cell cycle is classically considered as composed of 4 phases named G_1 (gap 1), S (DNA synthesis), G_2 (gap 2) and M (mitosis). One challenge of our modelling study was to determine the expression of the parameters d_i and $K_{i \rightarrow i+1}$ mentioned in the model (8.1) for each phase of the cell cycle ($i = 1 \dots 4$). We concentrated our efforts on $K_{i \rightarrow i+1}$, assuming that $d_i = 0$ for all $i = 1 \dots 4$. Note that if one assumes a constant death rate d for all phases, then it is included (negatively) in the growth exponent λ , which in fact may be thought of as the minimal added artificial death rate that stabilizes the cell population, as results from equations (8.2).

To get an expression for these transition rates, we used (8.4) to reduce the identification problem to the determination of the distribution of the duration of the phases of the cell cycle within a cell population. Other authors [STH⁺08] have used comparable modelling to investigate the cell cycle in cell populations, but the novelty of our contribution in this section is that we have used recent image data on individual cells that enabled us to assess the variability of cell cycle phase durations in populations of cells.

FUCCI is the acronym of fluorescent ubiquitination-based cell cycle indicator. This is a recently developed technique that allows tracking progression within the cell cycle of an individual cell with a high degree of contrast [SSKM⁺08, SSOH⁺08]. The FUCCI method consists in developing two fluorescent probes indicating whether a tracked cell is in the G_1 phase or in one of the phases S , G_2 or M of the cell cycle. The authors fused red- and green-emitting fluorescent proteins to proteins called Cdt1 and Geminin. Cdt1 and Geminin oscillate reciprocally: Cdt1 level is highest in the G_1 phase and falls down when the cell enters the S phase, whereas Geminin level is highest in the S , G_2 and M phases and falls when the cell enters the G_1 phase. Let us mention that Cdt1 and Geminin are degraded due to the process of ubiquitination, which is what is referred to (“U”) in the name of the reporter method. Consequently, the nucleus of a FUCCI cell fluoresces in red when this cell is in the phase G_1 , and in green when it is in S , G_2 or M phases.

This method allows to measure the time a tracked cell spends in the G_1 phase and the remaining part $S/G_2/M$ of the cell cycle. By tracking each cell in a population (note that cell tracking is not a completely trivial imagery problem because in liquid media cells move) we can get the distributions of the duration of these phases within the population, and so we can deduce the probability density functions of the random variables representing the duration of these phases (see below Subsection 8.5.3 for details).

8.5.2 Analysis of the experimental data

We used for the parameter identification procedure FUCCI data transmitted to us within the C5Sys EU project by G. van der Horst’s team, Erasmus University Medical Center, Rotterdam, The Netherlands. The cell lines were obtained by S. Saito at Erasmus University by recloning cell cycle phase markers (in the proper color combination, see further) and generating/analyzing NIH 3T3 cells (mouse embryonic fibroblasts) proliferating in a liquid medium. The data processed in the identification procedure thus consisted of time series of intensities

recording the red and green fluorescences emitted by individual NIH 3T3 cells proliferating within an *in vitro* homogeneous population. These cells had not been preliminarily synchronized, which means that they were initially at different stages of the cell cycle. The intensities had been recorded every fifteen minutes, over approximately 38 hours. A graph representing such a time series is presented on Figure 8.1.




Figure 8.1: Example of a time series of the intensity of red (deep gray) and green (light gray) fluorescences obtained by using the FUCCI method on a NIH 3T3 cell within a population in liquid medium.

We considered in the mass of data (about 2000 tracked cells) that were available to us only those (about 50) with at least the duration of a complete cell cycle, and measured the duration of the G_1 and $S/G_2/M$ phases within this cell cycle (note here that since all the cells that were kept for parameter identification were alive from the beginning of the experiment until its end, the assumption of a zero death rate in the model is in full accordance with these particular experimental conditions). The end of a cell cycle is characterised by a fast disappearance of the green fluorescence, so that it was not difficult to measure the duration of the cell cycle on our data. During the transition from G_1 to S , red and green fluorescences overlap, so that it is not so easy to determine the duration of phase G_1 . In agreement with our biologist partners, we decided to define the end of phase G_1 as the time at which red fluorescence was maximum before decreasing. The duration of phase $S/G_2/M$ was obtained by subtracting the duration of phase G_1 from the duration of the cell cycle. This method is summarized on Figure 8.2.

8.5.3 Expression of the transition rates without circadian control

With these processed data, we obtained 55 figures on individual cells for the duration of the cell cycle, divided in G_1 and $S/G_2/M$ phases. The mean value of the duration of the cell cycle was about 17.1 h (s.d.: 4.5 h), the one of G_1 was about 7.2 h (s.d.: 2.7 h), and thus the one of $S/G_2/M$ was about 9.8 h (s.d.: 3.0 h).

We rounded each duration to the nearest hour. The distributions of the durations of G_1 and of $S/G_2/M$ within the population were fitted to experimental data by using Gamma laws. The corresponding curves are presented on Figure 8.3. We tested several models to fit

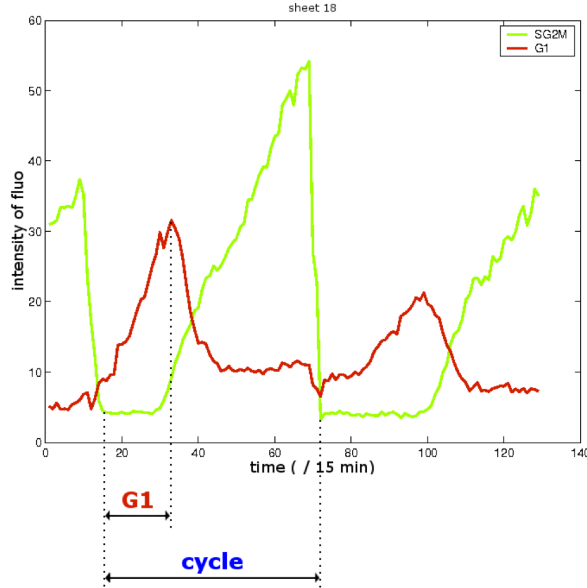


Figure 8.2: Graphic method used to determine the duration of the cell cycle and the one of G_1 phase. The duration of phase $S/G_2/M$ was deduced by subtracting the duration of phase G_1 from the duration of the cell cycle.

experimental data. We excluded laws that had support not strictly contained in \mathbb{R}_+ , such as normal laws; similarly, we excluded laws that assumed a maximum age in phase, which is impossible to define naturally and furthermore results in difficulties when identifying their parameters. We chose Gamma laws because they allowed a good (phenomenological) fit to our experimental data while keeping a reasonable number of parameters to be estimated.

Moreover, there is a clear physiological basis to this choice of the Gamma distribution: recall that, if the parameter α is an integer, it is the law, often used to represent probabilities of waiting times of the sum of α i.i.d. random variables representing waiting times (here within G_1 or $S/G_2/M$ of times between triggerings of crucial switches in a cascade of protein expressions leading to a phase transition, *e.g.*, G_1/S), each one of them following an exponential law with the same parameter β . Such an explanation, or parts of it, has been proposed in this context or others dealing with gene or protein expression by many authors, let us only mention [CKM⁺09, MA97, SS08]. Note that here the Gamma distribution is not used, as in [BPMM03] and references therein, to represent the distribution of maturation times for cells performing a fixed number of divisions, but the distribution of times spent in a phase of the cell division cycle, supposed to be constituted of a cascade of (unidentified, hidden physiological) switches, as mentioned above, hence its as phenomenological as physiological justification in our case.

For all $x \geq 0$, we thus used the following probability density functions, where Γ is the Gamma function:

$$\varphi_i(x) = \frac{1}{\Gamma(\alpha_i)} (x - \gamma_i)^{\alpha_i - 1} \beta_i^{\alpha_i} e^{-\beta_i(x - \gamma_i)} \mathbb{1}_{[\gamma_i; +\infty[}(x) \quad i = 1, 2, \quad \text{where}$$

$\alpha_1 = 8.28$, $\beta_1 = 1.052h^{-1}$, $\gamma_1 = 0h$, $\alpha_2 = 3.42$, $\beta_2 = 1.47h^{-1}$, $\gamma_2 = 7.75h$ and $\mathbb{1}_{[\gamma_i; +\infty[}$ is the indicator function of interval $[\gamma_i; +\infty[$. These parameters led to a mean duration and a

standard deviation on \mathbb{R}_+ respectively of $7.9h$ and $2.7h$ for the G_1 phase and of $10.1h$ and $1.3h$ for the $S/G_2/M$ phase. These figures are very close to the ones related to the raw experimental data mentioned above. The main difference resides in the $S/G_2/M$ data and is due to the fact that we have identified a high position parameter ($\gamma_2 = 7.75h$) for the Gamma distribution in $S/G_2/M$, which may be interpreted as an “incompressible” minimum duration for the $S/G_2/M$ part of the cell division cycle in our observed cell population, hence, *en passant*, an indirect measure of this cell physiological parameter, which is the minimum age a cell has to spend in $S/G_2/M$ before being able to process further and divide. Similarly, the figure 8.28 for α_1 , compared to 3.42 for α_2 may be interpreted - a speculation- as due to the presence of many more biological switches (and resulting stopping times) in G_1 than in $S/G_2/M$.

Note that the FUCCI technology only enables us to distinguish between cells in G_1 and $S/G_2/M$, without distinction between S , G_2 and M . However, the method used to identify phase transitions relies in fact on the probability distribution of durations of phases. Since the duration of the phase M is known to be most of the time very short, with almost zero variability within cell populations, it would be legitimate to consider it as fixed, as 1 hour, say, and that the recorded variability of $S/G_2/M$ is in fact the variability of S/G_2 . Thus, we shall consider that we were dealing in this identification process with a transition function from S/G_2 to M instead of the one from $S/G_2/M$ to G_1 . Under these assumptions, we will apply our optimization problem to an age-structured model accounting for 3 phases of the cell cycle, G_1 , S/G_2 and M . Then we can see that we have, thanks to FUCCI reporters, accessed the main two checkpoints, G_1/S and G_2/M .

Moreover, we have seen that the joint phase $S/G_2/M$ has an incompressible minimum duration γ_2 . This allows us to give a constant time of $\gamma_3 = 1h$ for the M phase, that we deduce from γ_2 . This constant time can be modeled by a transition rate from phase M to phase G_1 equal to 0 up to γ_3 and equal to a large positive value afterwards.

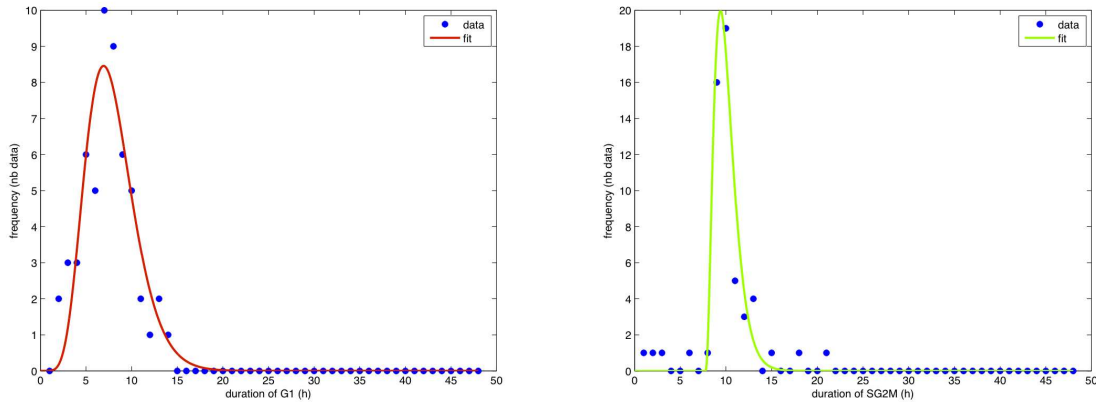


Figure 8.3: Gamma laws (solid line) (multiplied by a coefficient equal to the total number of data, *i.e.* 55) that fit experimental data (dots) for the distribution of the duration of G_1 (left) and of $S/G_2/M$ (right).

Let us also mention that in a close, but different context (linear age-structured population, one proliferation phase and presence of a death term), others have proposed an inverse problem method to identify, without any assumption on a probabilistic model for the cell cycle time,

the birth and death functions [GOP02].

Let us now come back to the 3-phase mathematical model ($I = 3$ in model (8.1)).

As the experimental data were performed *in vitro* in a liquid medium, with no intercellular communication, and as cells had not been synchronized prior to the experiment, we can consider that there was no time-dependent control whatsoever on the growth process at the cell population level. We thus assumed that the transition rates from G_1 to S/G_2 ($K_{1 \rightarrow 2}$), from S/G_2 to M ($K_{2 \rightarrow 3}$) and from M to G_1 ($K_{3 \rightarrow 1}$) did not depend on time, but only on the age of cells in the two phases. From the expression of the cumulative distribution function mentioned in Section 8.3:

$$\int_0^x \varphi_i(\xi) d\xi = 1 - e^{-\int_0^x K_{i \rightarrow i+1}(\xi) d\xi} \quad i = 1, 2,$$

we deduce:

$$K_{i \rightarrow i+1}(x) = \frac{\varphi_i(x)}{1 - \int_0^x \varphi_i(\xi) d\xi} \quad i = 1, 2 \tag{8.7}$$

where φ_i represents the experimentally determined probability density function of the random variable representing the age duration in phase i .

The graphs of the transition rates we obtained from formula (8.7) and experimental data are presented on Figure 8.4.




Figure 8.4: Transition rates from G_1 to S/G_2 (left) and from S/G_2 to M (right). These rates are functions of age of cells in the phases only.

8.5.4 Model simulations

As we are interested in the role of circadian rhythms on cell cycle progression, we consider transition rates of the form : $K_{i \rightarrow i+1}(x, t) = \kappa_i(x) \cdot \psi_i(t)$ ($i = 1, 2, 3$), where κ_i corresponds to the transition rates we identified in Subsection 8.5.3, and ψ_i represents a 24h-periodic control exerted by circadian rhythms on cell cycle progression at phase transitions.

8.5.4.a Internal validation To make sure that our numerical results were in agreement with the biological data that we used to build our model (“internal validation”), we performed simulations in the case of no time control, that is $K_{i \rightarrow i+1}(x, t) = \kappa_i(x)$ ($i = 1, 2, 3$), where the

$\kappa_i(x)$ were given by the expression in (8.7) for $i = 1, 2$ and $\kappa_3(x)$ equal 0 for $x \leq \gamma_3$ and a large positive value (we took 1000) for $x \geq \gamma_3$ ($\kappa_i(x) = K_{i \rightarrow i+1}(x)$). Figure 8.5 presents the time evolution of the percentage of cells in phases G_1 and the sum of cells in phases S/G_2 and M over the duration of one cell cycle resulting from numerical and biological experiments (biological data were preliminarily synchronized “by hand”, *i.e.*, by deciding that all cells were at age zero at the beginning of simulations). We had to reduce this comparison to the duration of one cell cycle because we had not enough biological data to represent the whole population. We can nevertheless notice that modeled numerical data were very close to raw biological data. We are thus entitled to conclude that the model and the method we have used to represent the proliferation phenomenon and fit our experimental data may have led us close to biological likelihood.




Figure 8.5: Time evolution of the percentages of cells in G_1 (red or deep gray) and the sum of cells in S/G_2 and M (green or light gray) phases from biological data (dashed line) and from numerical simulations (solid line). Our model results in a good approximation of the biological data.

8.5.4.b Numerical simulations In the case of no time control ($\psi_i \equiv 1$, $i = 1, 2, 3$), we studied the time evolution of the percentages of cells in G_1 , S/G_2 and M phases of the cell cycle and the evolution of the total density of cells. These results are presented on Figure 8.6. We can notice that oscillations are damped and that the percentages rapidly reach a steady state. This phenomenon, which has long been known in cell population dynamics [Ari95, AK93, AS97, CMTU01, MD86] as asynchronous growth, is the result of desynchronization of cells through the cell cycle: although the cells were taken initially all in phase G_1 with age 0, the variability in the duration of phases G_1 and S/G_2 that we described above induces some variability in the position of cells through their cycling. In our experimental case, as mentioned above, the exponential growth exponent λ computed by using Lotka’s equation (8.3) was equal to $0.039h^{-1}$.

Then, we introduced a circadian control modeled by functions ψ_i ($i = 1, 2, 3$) in the transition rates. For each phase, G_1 and S/G_2 and M we took continuous piecewise 24h-




Figure 8.6: Time evolution of the percentages of cells G_1 (left, red or deep gray), S/G_2 (left, green or light gray) and M (left, dotted blue) phases and of the total density of cells (right) evolving free of any time control. Oscillations are damped due to cell desynchronization. Population growth tends to be exponential with a rate equal to $0.039h^{-1}$.

periodic cosine functions, represented on Figure 8.7, defined on $[0;24]$ by:

$$\psi_1(t) = \cos^2(2\pi(t - 16)/12)\mathbb{1}_{[13;19]}(t) + \varepsilon, \quad \psi_2(t) = \cos^2(2\pi(t - 3)/12)\mathbb{1}_{[0;6]}(t) + \varepsilon, \quad \psi_3(t) = 1 \quad (8.8)$$

where $\varepsilon = 10^{-10}$ ensures $\psi_i > 0$ (which may be shown sufficient to imply irreducibility of matrix \mathbb{M} and thus applicability of the Perron-Frobenius theorem).

These phenomenological cosine-like functions standing here for physiological circadian control on cell cycle phase transitions in fact represent the local impact of the central circadian clock control onto the cell division cycle. This impact may be thought of either as retransmitted by a direct action of glucocorticoids - known to be synthesized on a circadian basis by the corticosurrenal gland [BK01] - directly on the G_1/S transition via stimulation of p27 and inhibition of cMyc [AOB⁺04, BCM⁺98], or as retransmitted by another relay involving local circadian clocks and control by Bmal1 of the complex Cyclin B-Cdk1 that controls the G_2/M transition [MYM⁺03]. In the absence of an actual experimental identification of the main gating variables (Cyclin E-Cdk2 on G_1/S and Cyclin B-Cdk1 on G_2/M) and the circadian control on them, we have chosen to represent them by such truncated cosines.

The 12h delay between the definition of ψ_1 and ψ_2 is suggested by biological observations, that teach us that circadian controls are exerted on the main checkpoints G_1/S and G_2/M and that proteins $p21$ and $Wee1$, known circadian controls on these checkpoints, are expressed in antiphase [GCRG⁺08]. Furthermore, and completely independently of such physiological knowledge, we remarked by varying the phase delay that this value of 12 hours maximized the population growth exponent λ . We thus obtained periodic transition rates $K_{i \rightarrow i+1}(x, t) = \psi_i(t) \cdot \kappa_i(x)$ of mixed origin, estimated by cosines ($\psi_i(t)$) and experimentally determined ($\kappa_i(x)$).

In this case, we observed that oscillations of the percentages of cells in G_1 , in S/G_2 and in M were no more damped (see Figure 8.8). In fact, because of circadian entrainment, cells were more synchronized in the cell cycle and consequently divided approximately all together. Moreover, the exponential growth rate was about $0.024h^{-1}$, which means that the population




Figure 8.7: Graphic definition of the 24h-periodic functions ψ_1 (left) and ψ_2 (right) modelling circadian control on the cell cycle. ψ_1 controls G_1 to S transition, and ψ_2 controls G_2 to M transition.

dynamics was slower than the one without time control we investigated above. Thus, in this case, circadian rhythms allowed cells to be more synchronized and to divide slower.




Figure 8.8: Time evolution of the percentages of cells in phases G_1 (left, red or deep gray), S/G_2 (left, green or light gray) and M (left, dotted blue) and of the total density of cells (right) submitted to circadian control. Oscillations are not damped any more. Population growth tends to be exponential with a rate equal to $0.024h^{-1}$.

These first numerical simulations tend to be in agreement with the biological hypothesis according to which populations of cells that can escape circadian control, such as cancer cell population could do, proliferate faster than populations of cells submitted to circadian entrainment.

As shown in [CGL09, CGL11, CGP07, CMP06, CMP07], there cannot be a general theoretical result for periodic control on cell cycle transitions, comparing proliferation in controlled and uncontrolled situations. Nevertheless, it could be that the particular form of the periodic control investigated here, *i.e.*, the product of an age-dependent component which may be qualified as hazard rate $\kappa_i(x)$ of a Gamma distribution for phase duration, and of a periodic time-dependent component $\psi_i(t)$ with optimized phase shift between the two transition control functions ψ_1 and ψ_2 , does result in slowing down the cell cycle speed. This remains to

be both experimentally and theoretically investigated.

8.5.4.c Modelling cancer cells We consider two cell populations called cancer cells and healthy cells. We make them differ only by the circadian time controls ψ_i between cell cycle phases i and $i + 1$, and we assume that there is no interaction between the two populations, healthy and cancer. We took for this circadian control a continuous truncated piecewise cosine function (*i.e.*, \cos^2 times an indicator function) for each phase.

For healthy cells, we took the functions described in Section 8.5.4, that is, we located the circadian control around 3 a.m. for the transition from S/G_2 to M (Figure 8.10) and around 4 p.m. for the transition from G_1 to S/G_2 . We assumed that cancer cell populations still obey circadian control at these main checkpoints but, like in [ALG07b], we modeled their behavior by a looser answer to the signal (Figure 8.9):

$$\begin{aligned}\psi_1^C(t) &= \frac{5}{8} \cos^2(2\pi(t - 16)/12) \mathbb{1}_{[13;16]}(t) + \frac{5}{8} \cos^2(2\pi(t - 16)/3) \mathbb{1}_{[16;24] \cup [0;4]}(t) + \varepsilon, \\ \psi_2^C(t) &= \frac{5}{8} \cos^2(2\pi(t - 3)/12) \mathbb{1}_{[0;3]}(t) + \frac{5}{8} \cos^2(2\pi(t - 3)/3) \mathbb{1}_{[3;15]}(t) + \varepsilon, \\ \psi_3^C(t) &= 1\end{aligned}$$

Then the cancer cells can go from phase G_1 to phase S/G_2 between 1 p.m. and 4 a.m. (15 hours) while healthy cells can do it between 1 p.m. and 7 p.m. (6 hours) We assumed that the counterpart of this looser answer to the signal is a lower maximal value for ψ_i^C than for ψ_i defined in (8.8). We set this maximal value such that $\int_0^T \psi_i^C(t) dt = \int_0^T \psi_i(t) dt$. Thanks to this looser answer to the circadian control, cancer cells get a proliferating advantage against healthy cells (Figure 8.12 below): the growth rate of healthy cells is $0.024h^{-1}$ whereas the growth rate of cancer cells is $0.026h^{-1}$.




Figure 8.9: Graphic definition of the 24h-periodic functions ψ_1^C (left) and ψ_2^C (right) modelling the looser circadian control on the cell cycle of cancer cells. ψ_1^C controls G_1 to S transition, and ψ_2^C controls G_2 to M transition.

8.6 A numerical optimization problem with toxicity constraints

In the previous sections we have described the effect of drugs in chemotherapy, an age-structured cell population dynamic model that can take into account phase-specific drugs

like 5-FU, the objectives and constraints considered in chemotherapy optimization problems that have been previously published. Those three topics can be seen as components of an optimization model. In order to get quantitative results, we identified the parameters of the model we had chosen in Section 8.5. Then, one has to choose an optimization procedure to solve the optimization problem considered.

When choosing an optimization procedure, one first needs to identify what are the optimization variables. For chemotherapy optimization, there are two main situations: either the optimization variables are some parameters of a predefined infusion scheme [Web90, PA95, ALG09] or they are the infusion scheme itself, represented by a time-dependent control function $g(t)$ as in our case (Section 8.4.1).

When the information provided by Pontryagin's maximum principle is enough to know the optimal control, as in [dLMS09, FP00, KS06, LMS11], it gives the control. The shooting method was also used in [LMS11] for chemotherapy optimization. Ledzewicz et al. considered two drugs that act on a Gompertzian model [Lai64].

Another alternative is to use a direct method. Direct methods consist of a total discretization of the control problem and then of solving the finite dimensional optimization problem obtained. The discretization of an optimal control problem results in an optimization problem with a large number of variables. The theory of differentiable optimization is the classical tool for such problems [Ber95, BGLS06, NW99]. This approach has been chosen in [PPdS95] with a gradient algorithm and in [BCL06] with a Uzawa algorithm. This is the approach we choose in this section.

In order to find global optima that can not be found in general by differentiable optimization, some authors use stochastic algorithms to solve the discretized problem, like for instance simulated annealing [AHL06] and CMA-ES [VOA10].

8.6.1 Optimization of the Floquet eigenvalue

In this work, we want to solve the problem of minimizing the asymptotic growth rate of the cancer cell population while keeping the asymptotic growth rate of the healthy cell population over a prescribed threshold. We modeled the cell population dynamics by a McKendrick model physiologically controlled by a circadian clock, considering a phase-dependent drug acting on transitions (8.1). We obtained the optimization problem (8.6) and we identified the main parameters.

For the numerical resolution of this problem, we firstly discretize the problem (8.1) with the discretization scheme presented in Section 8.3.4. Then, we get the problem of optimizing the first eigenvalue with control in a set of matrices. That is, we study the optimization of the growth rate in the discretized model. We chose a discretization step of 6 minutes because it may be considered as a lower limit to the half-life time of 5-FU in the plasma [BDP⁺00, DH89, PLK⁺93], which is most likely even lower than the half-time of its downstream molecular effects at the cell level, our concern here. The oldest ages represented in the discretization scheme are 10 days for each phase except for mitosis (phase M) where we chose 2h.

Unlike in [CGL09, CGL11], we will refer to the Perron eigenvalue when we consider non-negative square matrices and to the Floquet eigenvalue when we study the growth exponent of a time-periodic controlled population. Indeed, if we denote by ρ the Perron eigenvalue of the matrix M defined by the discretization scheme in Section 8.3.5, which means that $Mu = \rho u$ and ρ has maximal modulus, then the Floquet eigenvalue, λ can then be approximated by $\frac{1}{T} \log(\rho)$.

By the Perron-Frobenius theorem [BP94], we know that if M is nonnegative and irreducible, its principal eigenvalue $\rho(M)$ is positive and is a simple eigenvalue. Moreover, as mentioned above, the principal eigenvector is unique up to a normalization and can be chosen such that $u(M) \geq 0$. In our setting (irreducibility of the matrix and the Perron-Frobenius theorem), one can naturally define a function ρ from the set of nonnegative and irreducible real matrices in \mathbb{R}^n into \mathbb{R}_+ , that to a matrix associates its principal eigenvalue.

General eigenvalue optimization of non symmetric matrices is a difficult (non convex, non differentiable) problem: see [LO96] and [OW88] for two algorithms dealing with this problem. However, for positive matrices, as the principal eigenvalue is simple, this implies that ρ is differentiable. Indeed, by Proposition 7.6 [Kat66], denoting v and u the left and right eigenvectors of a matrix M associated to a simple eigenvalue ρ , the derivative of ρ at M can be written as:

$$\frac{\partial \rho}{\partial M_{ij}} = v_i u_j$$

Thus, as the objective function is differentiable, the theory of differentiable optimization applies.

As stressed by Overton in [Ove91], there are various possibilities for the computation of the eigenvalue and eigenvectors. Here, we consider sparse nonnegative matrices with a simple principal eigenvalue: the power method applies and, unlike direct methods or inverse iterations, it only needs matrix-vector products, which is valuable with a large sparse matrix.

The Perron (or Floquet) eigenvalue optimization problem with explicit constraints can be written as:

$$\min_{M \in h(C)} f(\rho(M)) \quad (8.9)$$

We assume that f is a real-valued twice continuously differentiable function; C a compact convex set and we denote P_C the orthogonal projection on C ; h is a twice continuously differentiable function such that $\mathcal{M} = h(C)$ is a set of nonnegative irreducible matrices. We have shown in Proposition 7.1 that this problem is NP-hard in general.

We may also need implicit constraints on the eigenvalues. The Perron (or Floquet) eigenvalue optimization problem with $K = |I| + |J|$ implicit constraints can be written as:

$$\begin{aligned} \min_{x \in C} f_0(\rho(h_0(x))) \\ f_k(\rho(h_k(x))) = 0, \quad k \in I \\ f_k(\rho(h_k(x))) \leq 0, \quad k \in J \end{aligned} \quad (8.10)$$

To solve this non convex problem, we use the method of multipliers [Ber82], which solves a sequence of non constrained optimization problems (8.9) whose solutions converges to the solution of the constrained problem (8.10). Let $F_k(x) = f_k(\rho(h_k(x)))$. Given $c > 0$, we call augmented Lagrangian of the Problem (8.10) the function L_c defined by

$$L_c(x, \mu) = F_0(x) + \sum_{j \in I} (\mu_j F_j(x) + \frac{c}{2} F_j(x)^2) + \frac{1}{2c} \sum_{j \in J} (\max(0, \mu_j + c F_j(x))^2 - \mu_j^2)$$

The method of multipliers consists in the following scheme, starting with μ^0 :

$$\begin{aligned} x^k &= \arg \min_x L_{c_k}(x, \mu^k) \\ \mu_j^{k+1} &= \mu_j^k + c_k F_j(x^k), \quad j \in I \\ \mu_j^{k+1} &= \max(0, \mu_j^k + c_k F_j(x^k)), \quad j \in J \end{aligned} \quad (8.11)$$

where the minimization is understood to be local in a ball within which x is the unique local minimum of Problem (8.10). Under classical assumptions (see [Ber82] for instance), the method of multipliers converges to a stationary point of the constrained optimization Problem (8.10).

For the resolution of the non constrained problems (8.9), we used the coupled power and gradient iterations algorithm developed in Section 7.4. It was designed for Perron vector optimization but by replacing the auxiliary vector w by the left eigenvector v in all the expressions, we can easily adapt the coupled power and gradient iterations to the Perron value optimization problem. This comes from the similarity of the expressions of the gradient for the two problems stated respectively in Propositions 7.6 and 7.7.

In our setting, at each time, the control $g_i(t)$ is the effect at the cell level of the drug infusion at time t on the transition rate from phase i to phase $i + 1$. No drug corresponds to $g_i(t) = 0$, a transition-blocking infusion corresponds to $g_i(t) = 1$. If the modeled drug is 5-FU, it acts on phase S (and thus on the aggregated phase S/G_2) on the DNA, resulting in damaged DNA and subsequent blocking control at the G_2/M transition only; we then have $g_1(t) = g_3(t) = 0$ for all t , and g_2 only is controlled. The discretized control x will be the array of the infusion time step by time step and drug by drug. As we search for 24h-periodic controls with one drug, we only need to define x on one day, i.e. $x \in \mathbb{R}^{N_T}$, where $N_T = T/\Delta t$ is the number of time steps. The requirement that $0 \leq g_2(t) \leq 1$ corresponds to $x \in C = [0, 1]^{N_T}$.

The various functions h_k , one by type of cell considered (healthy or cancer, bone cell or intestine cell...), represent the dependence of the model upon drug infusion. Given a discretized infusion strategy x we build the matrices $h_k(x)$, that are the nonnegative matrices modeling the discretized dynamics of each cell population under the drug infusion x . Transitions from one phase to the other are described by the transition rates $K_{i \rightarrow i+1}(t, x)$. As we take them with the form $K_{i \rightarrow i+1}(t, x) = \kappa_i(x) \psi_i(t) (1 - g_i(t))$ where $\kappa_i(x)$ is the transition rate of the cell without circadian control identified in Section 8.5.1, $\psi_i(t)$ is the natural circadian control and $g_i(t)$ is the effect of the drug on the transition from phase i to phase $i + 1$, we have $K_{i \rightarrow i+1}^{k,j} = \kappa_i^j \psi_i^k (1 - x^k)$. The Perron value of $h_k(x)$ is denoted by $\rho(h_k(x))$.

The objective will be the minimization of the growth rate (for cancer cells) $f_0(r) = \frac{1}{T} \log(r)$. For the constraints, we will consider a lower bound Λ (toxicity threshold) for the growth rate of healthy cells by $f_1(r) = -\frac{1}{T} \log(r) + \Lambda$.

8.6.2 Determination of the gradient of the objective function

In this section, we give an analytical expression of the gradient of the objective, that can be computed efficiently provided the left and right Perron vectors of the matrix \mathbb{M} defining the discretized dynamics (Section 8.3.5) are given.

From Proposition 7.6, we know that for all $i, j \leq I(j_{max} + 1)$,

$$\frac{\partial \rho}{\partial \mathbb{M}_{i,j}} = u_i v_j .$$

As \mathbb{M} is defined by

$$\mathbb{M} = M^{N_T} \dots M^2 M^1 ,$$

we introduce U^t the number of cells of each age at timestep t starting from $U^0 = u$ and V^t the discretized adjoint vector:

$$\begin{aligned} U^t &= M^{t-1} \dots M^1 u \\ V^t &= v M^{N_T} M^{N_T-1} \dots M^{t+1} . \end{aligned}$$

Then, we have

$$\frac{\partial \rho}{\partial M_{i,j}^t} = U_i^t V_j^t$$

From the formulas in Section 8.3.5, if we denote $X_k = \{i \mid (k-1)(j_{max}+1) < i \leq k(j_{max}+1)\}$ and $X_k(1) = (k-1)(j_{max}+1) + 1$, we have

$$M_{i,j}^t = \begin{cases} \beta_i^t & \text{if } i = j + 1, \\ \gamma_j^t & \text{if } i = X_{k+1}(1) \end{cases}$$

with the convention that $I+1 = 1$. We have $\beta_i^t = Ak_{i-X_k(1)+2, i-X_k(1)+1}^t$ for $k = 1 \dots I$ and $i \in X_k, i \neq k(j_{max}+1)$ and $\gamma_j^t = Bk_{1, j-X_k(1)+1}^t$ for $k = 1 \dots I$ and $j \in X_k$. The block matrices Ak^t and Bk^t are defined in Section 8.3.5. Indeed, β_i^t only depends on $K_{k \rightarrow k+1}^{t+1, i}$ and $d_k^{t+1, i}$ for $i \in X_k$. Then,

$$\begin{aligned} \frac{\partial \rho}{\partial \beta_i^t} &= U_i^t V_{i-1}^t \\ \frac{\partial \rho}{\partial \gamma_j^t} &= U_{X_{k+1}(1)}^t V_j^t, \quad j \in X_k . \end{aligned}$$

By the chain rule, we obtain

$$\begin{aligned} \frac{\partial \rho}{\partial K_{k \rightarrow k+1}^{t+1, j}} &= \frac{\partial \rho}{\partial \gamma_j^t} \frac{\partial \gamma_j^t}{\partial K_{k \rightarrow k+1}^{t+1, j}} + \frac{\partial \rho}{\partial \beta_j^t} \frac{\partial \beta_j^t}{\partial K_{k \rightarrow k+1}^{t+1, j}} \\ \frac{\partial \rho}{\partial K_{k \rightarrow k+1}^{t, j}} &= U_{X_{k+1}(1)}^{t-1} V_j^{t-1} \frac{\partial \gamma_j^{t-1}}{\partial K_{k \rightarrow k+1}^{t, j}} + U_j^{t-1} V_{j-1}^{t-1} \frac{\partial \beta_j^{t-1}}{\partial K_{k \rightarrow k+1}^{t, j}} . \end{aligned}$$

We can follow further the chain rule if $K_k^{t, j}$ itself is indirectly controlled by drug infusions g_d^t .

$$\begin{aligned} \frac{\partial \rho}{\partial g_d^t} &= \sum_{k=1}^I \sum_{j \in X_k} \sum_{\tau=1}^{N_T} \frac{\partial \rho}{\partial \gamma_j^{\tau-1}} \frac{\partial \gamma_j^{\tau-1}}{\partial K_k^{\tau, j}} \frac{\partial K_k^{\tau, j}}{\partial g_d^t} + \sum_{k=1}^I \sum_{i \in X_k} \sum_{\tau=1}^{N_T} \frac{\partial \rho}{\partial \beta_i^{\tau-1}} \frac{\partial \beta_i^{\tau-1}}{\partial K_k^{\tau, i}} \frac{\partial K_k^{\tau, i}}{\partial g_d^t} \\ \frac{\partial \rho}{\partial g_d^t} &= \sum_{k=1}^I \sum_{j \in X_k} \sum_{\tau=1}^{N_T} U_{X_{k+1}(1)}^{\tau-1} V_j^{\tau-1} \frac{\partial \gamma_j^{\tau-1}}{\partial K_k^{\tau, j}} \frac{\partial K_k^{\tau, j}}{\partial g_d^t} + \sum_{k=1}^I \sum_{i \in X_k} \sum_{\tau=1}^{N_T} U_i^{\tau-1} V_{i-1}^{\tau-1} \frac{\partial \beta_i^{\tau-1}}{\partial K_k^{\tau, i}} \frac{\partial K_k^{\tau, i}}{\partial g_d^t} \end{aligned}$$

We assumed here that $K_k^{\tau, j}$ only depends on u^τ , so that the sum on τ is trivial. However, taking into account the PK-PD of the drug in a less trivial way would lead to $K_k^{\tau, j}$ depending on the whole drug infusion strategy.




Figure 8.10: *Drug and circadian controls, healthy cell population case.* Cosine-like functions modelling the drug and circadian controls for transition from G_1 to S/G_2 (dash-dotted line) and for transition from S/G_2 to M in healthy cells. The “natural” (drug-free) control for S/G_2 to M transition corresponds to the solid line, the drug-induced one to the dashed line.




Figure 8.11: *Drug and circadian controls, cancer cell population case.* Cosine-like functions modelling the drug and circadian controls for transition from G_1 to S/G_2 (dash-dotted line) and for transition from S/G_2 to M in cancer cells. The “natural” (drug-free) control for S/G_2 to M transition corresponds to the solid line, the drug-induced one to the dashed line.

8.6.3 Simulations

We considered two cell populations called cancer cells and healthy cells. In these simulations, we made them differ only by the circadian time controls ψ_i between cell cycle phases i and $i + 1$, and we assumed that there was no interaction between the two populations, healthy and cancer.

We took for this circadian control a continuous truncated piecewise cosine function (*i.e.*, \cos^2 times an indicator function) for each phase, as described in Section 8.5.4, that is, we located the circadian control around 3 a.m. for the transition from S/G_2 to M (Figure 8.10) and around 4 p.m. for the transition from G_1 to S/G_2 . We assumed that cancer cell populations still obey circadian control at these main checkpoints but, like in [ALG07b], we modeled their behavior by a looser answer to the signal (Figure 8.9). We assumed that the drug has the same effect on both populations, which couples their behaviors through the drug infusions.




Figure 8.12: Evolution of the population of cancer (blue, above) and healthy (green, beneath) cells without drug infusion during 12 days. We can see that the populations have different exponential growth rates ($\lambda_{cancer} = 0.026$ and $\lambda_{healthy} = 0.024$). Cancer cells proliferate faster than healthy cells.




Figure 8.13: Evolution of the population of cancer (blue, beneath) and healthy (green, above) cells with the drug infusion, starting at time 0, given by the algorithm. Healthy cells keep multiplying ($\lambda_{healthy} = 0.022$) while the cancer cell population is weakened ($\lambda_{cancer} = 0.019$). Contrarily to the case without drug, cancer cells proliferate slower than healthy cells.




Figure 8.14: Locally optimal drug infusion strategy (function g_2 , see text for details) found by the optimization algorithm.

Without drug infusion, the growth rate of cancer cells ($0.026h^{-1}$) is assumed to be larger than the one of healthy cells ($0.024h^{-1}$). This gives an evolution of the respective populations, cancer cells becoming more and more present: see Figure 8.12. After convergence of the method of multipliers, we get the locally optimal strategy, shown on Figure 8.14, defining on $[0; 24]$ the 24h-periodic function g_2 (recall that $g_1 = 0$ and $g_3 = 0$).

We can see the action of the locally optimal drug infusion strategy, provided by the optimization algorithm, on transition rates illustrated on Figures 8.10 and 8.11. This strategy restricts transition time durations from S/G_2 to M to three hours (between 1 am and 4 am) for both cell populations, whereas under the drug-free circadian control, transitions would be possible during 6 hours (between midnight and 6 am) for healthy cells and during 15 hours (between midnight and 3 pm) for cancer cells. Thus the physical meaning of this locally optimal infusion strategy is to forbid transitions from S/G_2 to M when cancer cells are under target while healthy cells are not (or very little), thus harming mostly cancer cells. Depending on the toxicity threshold Λ chosen, more aggressive drug infusions are possible.

By following the infusion strategy numerically determined by the optimization algorithm, we obtained that the growth rate of healthy cells was above the chosen toxicity threshold and that the growth rate of cancer cells was strongly weakened. This gave us a description of the evolution of the respective populations, which is illustrated on Figure 8.13.

We finally simulated the transition from the stationary state without drug to the stationary state with periodic drug infusion (Figure 8.15). After a transition of around 10 days, the treatment performs as expected ($\lambda_{cancer} = 0.019$ and $\lambda_{healthy} = 0.022$). We have thus pulled by this optimal infusion strategy the whole cell population from a state favorable to cancer cells to a state favorable to healthy cells.

8.7 Discussion and future prospects for this work

The problem of circadian control on the cell division cycle in cell populations and its possible applications in clinical oncology is a question of biological and clinical origin that has already been studied from a theoretical point of view in the mathematical setting of age-structured physiological cell population dynamics [CGL09, CGL11, CGP07, CLMP03, CMP06, CMP07]. Yet many unsolved questions remain, which may be due, in particular, to the scarcity of data




Figure 8.15: Daily mean growth rates for cancer (solid line) and healthy cells (dashed line) when starting drug infusions at time 0. After a 10-day transitional phase, the biological system stabilizes towards the expected asymptotic growth rate.

on parameters at the individual cell level in proliferating cell populations.

Taking advantage of quantitative measurements obtained by performing recent image analysis techniques of the cell division cycle in individual cells inside a population of non communicating proliferating cells of the same healthy lineage, we have focused in this work on studying age synchronization of cells with respect to cell cycle phases. Although these cells are far from an *in-vivo* situation, they give us precious indications on the age distribution within the cell division cycle in a homogeneous cell population without control exerted on its proliferation, neither by circadian clock messages, nor by applied drugs.

The drug-free experimental proliferation dynamics of this cell population is well approximated by Gamma distributions for cycle phase durations, for which we have shown that the growth exponent λ , first eigenvalue of the system, is increasing with the variabilities of these durations.

We assumed a multiplicative expression for both temporal controls, physiological (circadian) and pharmacological, onto cell cycle phase age-dependent transition kernels in the McKendrick model of cell proliferation with 3 phases. Our results on long-time drug-free cell population dynamics behavior, as shown on Figures 8.6 and 8.8, are consistent with the theoretical and experimental results presented in [CMTU01], that report classical asynchronous cell growth [Ari95, AK93, AS97, CMTU01, MD86, Per07], with theoretical works that report entrainment of the phase population densities by periodic control [Cla08], and comparable with the results presented in another modelling context in [AGLG11, ALG07b].

We have also been able to propose a new therapeutic optimization scheme under a toxicity constraint, controlling growth exponents in both cancer and healthy cell populations. We resolved this optimization problem by using the method of multipliers, which yielded a locally optimal drug infusion strategy.

We can see clear similarities between this infusion strategy and others found in the literature on cancer chronotherapeutic optimization, although a great variety of models may be considered (see [BCF12] for a comparison of these models). Indeed, in [ALG07b, BCL06] as well as in our case, the suggested infusion schedules set a maximal drug infusion flow when cancer cells get hurt by the drug while healthy cells do not. The modelling settings are however different with respect to the drug effect considered (on death rates in [ALG07b, BCL06, BCBLH10] or on proliferation rates, in the present study) and with respect to the model of cell population

dynamics chosen. In [BCL06], the solutions (drug delivery flows) to the optimization problem are much smoother than in the present study. This may be attributed to the fact that here, we represented control by blockade of proliferation (on cell cycle phase transition rates and the resulting growth coefficients), which is closer to a ‘tap open - tap closed’ problem than an action on death rates. Whereas most therapeutic control representations so far have been put on death rates (but note that both targets are considered in [KS06], and comparatively studied in [CGL11, CMP06, Lep09]), we know that anticancer drugs act mostly on proliferation, either by damaging the DNA, which results in subsequent cell cycle checkpoint (G_1/S or G_2/M) blockade via ATM and p53 in the case of cytotoxics, or by slowing down the G_1 phase in the case of cytostatics (growth factor inhibitors). Drugs that act directly on death rates, *e.g.*, by primarily enhancing the apoptotic cascade may exist (or will exist some day), but are not of common use in the clinic. Hence our choice of the drug target, transition rates, which looks more realistic to us than death rates.

Of note, the authors of [BCBLH10], also using an age-structured population dynamic model, but with delays, based their optimization procedure on the remark that the set of 24h-periodic strategies contains both best and worst strategies, depending on the deemphasizing between the position of the maximal effect of the circadian clock and of the drug infusion. Taking advantage of this remark, they aim at avoiding possible traps (*i.e.*, ‘pessimized infusion strategies’) when one has no precise idea about the optimal circadian time, thinking that it is advisable to propose a robust optimization approach by using a different period for the drug infusion scheme. Thus, they proposed drug infusion schedules that are not 24h-periodic, making their drug effects less dependent upon this dephasing (otherwise said, ‘shooting more safely in the dark’), which is a quite interesting point of view. In our case, since our approach is based on experimental data supposed to give us insight onto the target and on actual possibilities to reach it safely, we did not consider such a robust optimization approach.

Our optimization method relies on the assumption of differences between healthy and cancer cell population model parameters, namely drug-free time-dependent circadian control functions ψ_i on phase transitions. The clinical feasibility status of the proposed optimal drug infusion strategy is of course still questionable. Indeed, the ψ_i functions are thus far purely phenomenological (cosine-like functions) and, further, the local (tissue) drug effect proposed as control variable does not take into account tissue pharmacokinetics-pharmacodynamics (PK-PD) of any drug. Last but not least, in view of clinical applications, a whole-body physiologically based PK-PD model is still lacking. Nevertheless, by using this combination of physiologically based modelling of proliferation, mathematical analysis methods, cell imaging and statistical parameter identification techniques, and original optimization algorithms using eigenvalue control of growth processes, we propose the first steps of a rationale for therapeutic optimization in oncology at the molecular level in cell populations, healthy and tumour. We intend to complete these first steps in the future, as sketched in [Cla07, Cla09, Cla11], to get closer to the clinic.

Various measurements needed to identify parameters of our model control functions were still out of reach by the biological experiments performed in this pioneering study, and this is the reason why we used only a phenomenological representation (by plain cosines) of circadian control, but more measurements are expected to come from further experiments performed by biologists on samples of cell populations, healthy or cancerous, with or without circadian control.

In particular, in forthcoming recordings of FUCCI data on healthy and cancer cells, we will pay attention to parallel experimental measurements of the growth exponents (that are

inverses of doubling times multiplied by a factor $\ln 2$) in the proliferating cell populations at stake. We will thus identify experimentally the growth coefficient λ on cell population samples, in which we will simultaneously identify the parameters of our model from cell cycle phase duration distributions by FUCCI recordings, obtaining a theoretical growth exponent by solving the Euler-Lotka equation. This will allow us to validate (or falsify) the model. But thus far, such simultaneous measurements have not been made available to us.

In the immediate future, we intend to combine the McKendrick cell population model with a PK-PD model for 5-FU and Irinotecan, as the one presented in [LOD⁺10]. Hence, the optimization model would become by far more realistic since it would give an actual estimation of optimal drug flux instead of optimal drug effects. Moreover, the PK-PD model may give rise to phenomena that are not accessible with the present model.

CHAPTER 9

Conclusion

Throughout this thesis, we encountered the different steps of the resolution of an optimization problem. We proposed new models: the continuous PageRank optimization problem, a Markov decision process for spam detection, deformed HOTS algorithms for the ranking of web pages and a chemotherapy optimization model based on an age-structured cell population dynamics. Then, we developed algorithms for their resolution. We studied Markov decision processes with implicitly defined action spaces and provided efficient algorithms for the PageRank optimization problem. For the Perron eigenvalue and eigenvector optimization problems we gave an algorithm coupling power and gradient iterations, basing on the low rank of the matrix of the partial derivatives of the objective function and on the Master algorithm model of Polak. Each algorithm is associated with the corresponding proof of convergence. We implemented the algorithms in Scilab, Matlab or C language and we gave experimental results on small size, middle size and large size problems that demonstrate the scalability of the algorithms. We finally analyzed the optimal strategies found. In particular, we proved that there generally exists a threshold property for web ranking optimization problems.

An open question is related to the hardness of the global minimization of the Perron eigenvalue on polytopes. Actually, the polynomial time algorithms for Perron value optimization require that the controls are independent line by line. This gives a very raw relaxation in the Perron eigenvalue optimization problem that we studied here in the context of chemotherapeutics because the controls apply here on the subdiagonal of the matrix. What is more in our model, there is an implicit constraint that should be relaxed too. Does there exist polynomial time computable lower bounds with a better guarantee of quality? One could also

search for lower or upper bounds for each of the nonconvex Perron eigenvector or eigenvalue optimization problems considered here, each one being specific.

In this thesis, we have encountered several web ranking algorithms and we have studied how a webmaster could improve her ranking by manipulating the hyperlinks she controls. A natural question that arises is which algorithm a search engine should use. The answer will perhaps never be answered once and for all because the quality of a ranking is a subjective issue. However a web ranking algorithm should be scalable because of the size of the web, being linked to by more pages should increase the rank and the ranking should have some stability properties against perturbations and some resistance against web spamming. The Perron eigenvector optimization algorithm that we developed may help determining what kind of strategies will be chosen by spammers and more generally if well behaved webmasters will be penalized or not.

APPENDIX \mathcal{A}

An upper bound on the spectral gap of nonnegative matrices

In this appendix, we give an independent result on estimations of the spectral gap of a nonnegative matrix. The motivation for this study is that the size of the spectral gap gives the speed of convergence of the power method. As we have seen for Tomlin's HOTS algorithm in Chapter 6, a larger spectral gap is an argument in favor of an algorithm against an algorithm with a characteristic matrix that has a smaller spectral gap. We propose here a bound on the spectral gap based on matrix reversiblization and previous bound determined for symmetric matrices. Unfortunately, this bound may give no information, even for primitive matrices.

A.1 Reversiblization for nonnegative matrices

In [Fil91], Fill gives a technique that considers a reversible Markov chain associated to a non-reversible Markov chain such that they have the same stationary distribution. This technique, called reversiblization can be derived in additive and multiplicative reversiblization and we propose an extension to nonnegative (not necessarily stochastic) matrices.

We consider an irreducible nonnegative matrix M with left and right Perron vectors v and u and Perron root ρ . We investigate bounds on $|\lambda|$ for any other eigenvalue λ of A .

We define the reversed matrix \tilde{M} by

$$\tilde{M} = D^{-1}M^T D$$

where D is the diagonal matrix such that $D_{i,i} = \frac{v_i}{u_i}$. Note that \tilde{M} has the same spectrum as M and the same Perron vectors.

The multiplicative reversibilization of M is defined by

$$S(M) = D^{1/2}(M\tilde{M})D^{-1/2} = D^{1/2}MD^{-1}M^TD^{1/2}.$$

$S(M)$ is nonnegative, semi-definite, its Perron root is ρ^2 and one of its (left or right) Perron vector is $w = D^{1/2}u$ such that $w_i = \sqrt{u_i v_i}$. Note that even if M is irreducible, $S(M)$ may be reducible. This situation is not very problematic because a symmetric reducible matrix is easily tractable.

The additive reversibilization of M is defined by

$$A(M) = \frac{1}{2}D^{1/2}(M + \tilde{M})D^{-1/2} = \frac{1}{2}(D^{1/2}MD^{-1/2} + D^{-1/2}M^TD^{1/2}).$$

$A(M)$ is nonnegative and symmetric. Its Perron root and Perron vector are ρ and $w = D^{1/2}u$. It keeps the irreducibility property of M .

A.2 Evaluating the spectral gap of a nonnegative matrix

In [FG94], Friedland and Gurvits give a bound for the real part of the nonprincipal eigenvalues of a nonnegative matrix. It uses additive reversibilization and the result is the following:

$$\rho(M) - \Re(\lambda(M)) \geq \frac{1}{2}(\rho(M) - \max_{1 \leq i \leq n} m_{ii})\epsilon(M, u, v)^2$$

where

$$\epsilon(M, u, v) = \inf_{\emptyset \neq U \subset \{1, \dots, n\}, \text{card}(U) \leq \lfloor \frac{n}{2} \rfloor} \frac{\sum_{i \in U, j \in \{1, \dots, n\} \setminus U} m_{ij} v_i u_j + m_{ji} v_j u_i}{\sum_{i \in U} 2(\rho(M) - m_{ii}) v_i u_i}$$

It uses the fact that $\Re(\lambda(M)) \leq \lambda_2(A(M))$, the second eigenvalue of the symmetric matrix $A(M)$ and a bound for symmetric matrices [Fri92].

We propose a similar bound for the modulus of the eigenvalue instead of the real part. We get it thanks to multiplicative reversibilization.

Proposition A.1. *The spectral gap of a nonnegative matrix M is bounded by its multiplicative reversibilization's:*

$$|\lambda_2(M)|^2 \leq \sigma_2(S(M))$$

Proof. Let ρ^2 and w be the Perron eigenvalue and eigenvector of $S(M)$. Let σ_2 be its second eigenvalue (also singular value). Let λ_2, v_2 be a nonprincipal left eigenpair of M . We normalize v_2 such that $\|D^{-1/2}v_2\|_2 = 1$. As $w^T D^{-1/2}v_2 = u^T v_2 = 0$, we have:

$$\begin{aligned} \sigma_2 &= \max_{x \in \mathbb{C}^n, \|x\|_2=1, w^T x=0} x^* S(M)x \geq (D^{-1/2}v_2)^* S(M)(D^{-1/2}v_2) \\ &= v_2^* D^{-1/2} D^{1/2} M D^{-1} M^T D^{1/2} D^{-1/2} v_2 = (\bar{\lambda}_2 v_2^* D^{-1/2})(D^{-1/2} \lambda_2 v_2) = |\lambda_2|^2 \quad \square \end{aligned}$$

We can now use any bound on $\sigma_2(S(M))$ to get a bound on $|\lambda_2(M)|$.

Remark A.1. In case $S(M)$ is reducible, then one can show that $\sigma_2 = \rho^2$. This means that the bound is then useless.

Bibliography

- [AAD08] Tony Abou-Assaleh and Tapajyoti Das. Combating spamdexing: Incorporating heuristics in link-based ranking. In William Aiello, Andrei Broder, Jeanette Janssen, and Evangelos Milios, editors, *Algorithms and Models for the Web-Graph*, volume 4936 of *Lecture Notes in Computer Science*, pages 97–106. Springer Berlin / Heidelberg, 2008.
- [AD12] Marianne Akian and Sylvie Detournay. Multigrid methods for two-player zero-sum stochastic games. *Numerical Linear Algebra with Applications*, 19(2):313–342, 2012.
- [AG03] Marianne Akian and Stéphane Gaubert. Spectral theorem for convex monotone homogeneous maps, and ergodic control. *Nonlinear Analysis: Theory, Methods & Applications*, 52(2):637–679, 2003.
- [AGLG11] Atilla Altinok, Didier Gonze, Francis Lévi, and Albert Goldbeter. An automaton model for the cell cycle. *Interface focus*, 1:36–47, 2011.
- [AGN06] Marianne Akian, Stéphane Gaubert, and Laure Ninove. The T-PageRank: A model of self-validating effects of web surfing. In *2nd Multidisciplinary International Symposium on Positive Systems (POSTA 2006)*, volume 341 of *Lecture Notes in Control and Information Sciences*, pages 239–246. Springer-verlag Berlin, 2006.
- [AGN11] Marianne Akian, Stéphane Gaubert, and Roger Nussbaum. A collatz-wielandt characterization of the spectral radius of order-preserving homogeneous maps on cones. *arXiv:1112.5968*, 2011.
- [AGN12] Marianne Akian, Stéphane Gaubert, and Roger Nussbaum. Uniqueness of the fixed point of nonexpansive semidifferentiable maps. *arxiv:1201.1536*, 2012.
- [AH08] Jesse Alpert and Nissan Hajaj. We knew the web was big..., 2008. <http://googleblog.blogspot.fr/2008/07/we-knew-web-was-big.html#!/2008/07/we-knew-web-was-big.html>.

- [AHH⁺09] Jonathan Axtell, Lixing Han, Daniel Hershkowitz, Michael Neumann, and Nung-Sing Sze. Optimization of the spectral radius of a product for nonnegative matrices. *Linear Algebra and its Applications*, 430(5-6):1442–1451, 2009. Special Issue devoted to the 14th ILAS Conference.
- [AHL06] Zvia Agur, Refael Hassin, and Sigal Levy. Optimizing chemotherapy scheduling using local search heuristics. *Oper. Res.*, 54:829–846, 2006.
- [AK93] Ovide Arino and Marek Kimmel. Comparison of approaches to modeling of cell population dynamics. *SIAM J. Appl. Math.*, 53:1480–1504, 1993.
- [AL06] Konstantin Avrachenkov and Nelly Litvak. The Effect of New Links on Google PageRank. *Stochastic Models*, 22(2):319–331, 2006.
- [ALG07a] Atilla Altinok, Francis Lévi, and Albert Goldbeter. A cell cycle automaton model for probing circadian patterns of anticancer drug delivery. *Adv. Drug Deliv. Rev.*, 59:1036–1010, 2007.
- [ALG07b] Atilla Altinok, Francis Lévi, and Albert Goldbeter. A cell cycle automaton model for probing circadian patterns of anticancer drug delivery. *Adv. Drug Deliv. Rev.*, 59:1036–1010, 2007.
- [ALG09] Atilla Altinok, Francis Lévi, and Albert Goldbeter. Identifying mechanisms of chronotolerance and chronoefficacy for the anticancer drugs 5-fluorouracil and oxaliplatin by computational modeling. *Eur. J. Pharm. Sci.*, 36:20–38, 2009.
- [Alt99] Eitan Altman. *Constrained Markov decision processes*. Stochastic Modeling. Chapman & Hall/CRC, 1999.
- [ANTT02] Arvind Arasu, Jasmine Novak, Andrew Tomkins, and John Tomlin. Pagerank computation and the structure of the web: Experiments and algorithms. In *Proc. 11th International Conference on World Wide Web, Alternate Poster Tracks*, 2002.
- [AOB⁺04] M. J. Ausserlechner, P. Obexer, G. Böck, S. Geley, and R. Kofler. Cyclin D3 and c-MYC control glucocorticoid-induced cell cycle arrest but not apoptosis in lymphoblastic leukemia cells. *Cell death differ.*, 11:165–174, 2004.
- [Ari95] Ovide Arino. A survey of structured cell population dynamics. *Acta Biotheor.*, 43:3–25, 1995.
- [AS97] Ovide Arino and Eva Sanchez. A survey of cell population dynamics. *J. Theor. Med.*, 1:35–51, 1997.
- [Bal98] Egon Balas. Disjunctive programming: Properties of the convex hull of feasible points. *Discrete Applied Mathematics*, 89(1-3):3–44, 1998.
- [Bar51] M. S. Bartlett. An inverse matrix adjustment arising in discriminant analysis. *The Annals of Mathematical Statistics*, 22(1):107–111, 1951.
- [BBCP08] Fadia Bekkal Brikci, Jean Clairambault, and Benoît Perthame. Analysis of a molecular structured population model with polynomial growth for the cell cycle. *Mathematical and Computer Modelling*, 47:699–713, 2008.

- [BBCRB08] Fadia Bekkal Brikci, Jean Clairambault, Benjamin Ribba, and Benoît BPerthame. An age-and-cyclin-structured cell population model for healthy and tumoral tissues. *Journal of Mathematical Biology*, 57:91–110, 2008.
- [BCBLH10] Samuel Bernard, Branka Cajavec Bernard, Francis Lévi, and Hanspeter Herzel. Tumor growth rate determines the timing of optimal chronomodulated treatment schedules. *PLoS Computational Biology*, 6(3):e1000712, 2010.
- [BCF⁺11a] Frédérique Billy, Jean Clairambault, Olivier Fercoq, Stéphane Gaubert, Thomas Lepoutre, and Thomas Ouillon. Proliferation in cell population models with age structure. In *Proceedings of ICNAAM 2011, Kallithea Chalkidis (Greece)*, pages 1212–1215. American Institute of Physics, 2011.
- [BCF⁺11b] Frédérique Billy, Jean Clairambault, Olivier Fercoq, Stéphane Gaubert, Thomas Lepoutre, Thomas Ouillon, and Shoko Saito. Synchronisation and control of proliferation in cycling cell population models with age structure. *Mathematics and Computers in Simulation*, 2011. doi:10.1016/j.matcom.2012.03.005.
- [BCF12] Frédérique Billy, Jean Clairambault, and Olivier Fercoq. Optimisation of cancer drug treatments using cell population dynamics. In U. Ledzewicz, H. Schättler, A. Friedman, and E. Kashdan, editors, *Mathematical Methods and Models in Biomedicine*, pages 263–299. Springer, 2012. In press.
- [BCL06] Claude Basdevant, Jean Clairambault, and Francis Lévi. Optimisation of time-scheduled regimen for anti-cancer drug infusion. *Mathematical Modelling and Numerical Analysis*, 39:1069–1086, 2006.
- [BCM⁺98] N. Baghdassarian, R. Catallo, M. A. Mahly, P. Ffrench, F. Chizat, P. A. Bryon, and M. Ffrench. Glucocorticoids induce G1 as well as S-phase lengthening in normal human stimulated lymphocytes: differential effects on cell cycle regulatory proteins. *Exp. Cell Res.*, 240:263–273, 1998.
- [BCSU05] András Benczúr, Károly Csalogány, Tamás Sarlós, and Máté Uher. Spamrank – fully automatic link spam detection. In *1st International Workshop on Adversarial Information Retrieval on the Web*, 2005.
- [BCSV04] Paolo Boldi, Bruno Codenotti, Massimo Santini, and Sebastiano Vigna. Ubi-crawler: A scalable fully distributed web crawler. *Software: Practice & Experience*, 34(8):711–726, 2004. Webgraph available: <http://law.dsi.unimi.it>.
- [BDP⁺00] Guido Bocci, Romano Danesi, Antonello Di Paolo, Federico Innocenti, Giacomo Allegrini, Alfredo Falcone, Alessandro Melosi, Mario Battistoni, Gemma Barsanti, Pier Franco Conte, and Mario Del Tacca. Comparative pharmacokinetic analysis of 5-fluorouracil and its major metabolite 5-fluoro-5,6-dihydrouracil after conventional and reduced test dose in cancer patients. *Clin. Cancer Res.*, 6:3032–37, 2000.
- [Ber76] Dimitri P. Bertsekas. On the Goldstein - Levitin - Polyak gradient projection method. *IEEE transactions on automatic control*, AC-21(2):174–184, 1976.

- [Ber82] Dimitri P. Bertsekas. *Constrained Optimization and Lagrange multiplier method*. Academic Press, NY; republished by Athena Scientific, MA, 1997, 1982.
- [Ber95] Dimitri P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1995.
- [Ber05] Pavel Berkhin. A survey on pagerank computing. *Internet Mathematics*, 2:73–120, 2005.
- [BGH⁺04] Vincent D. Blondel, Anahí Gajardo, Maureen Heymans, Pierre Senellart, and Paul Van Dooren. A measure of similarity between graph vertices: Applications to synonym extraction and web searching. *SIAM Review*, 46(4):647–666, 2004.
- [BGLS06] J. Frédéric Bonnans, Jean-Charles Gilbert, Claude Lemaréchal, and Claudia Sagastizábal. *Numerical Optimization – Theoretical and Practical Aspects*. Universitext. Springer Verlag, Berlin, 2006.
- [BK01] R. M. Buijs and A. Kalsbeek. Hypothalamic integration of central and peripheral clocks. *Nature Rev. Neurosci.*, 2:521–526, 2001.
- [BL04] Anthony Bonato and Wilfrid Laurier. A survey of models of the web graph. In *Proc. of Combinatorial and Algorithmic Aspects of Networking*, pages 159–172. Springer, 2004.
- [BL06] Luca Becchetti and Stefano Leonardi. Using rank propagation and probabilistic counting for link-based spam detection. In *Proceedings of the Workshop on Web Mining and Web Usage Analysis (WebKDD)*. ACM Press, 2006.
- [BN05] Vincent D. Blondel and Yurii Nesterov. Computationally efficient approximations of the joint spectral radius. *SIAM J. Matrix Anal. Appl.*, 27(1):256–272, 2005.
- [BN10] Vincent D Blondel and Yurii Nesterov. Polynomial-time computation of the joint spectral radius for some sets of nonnegative matrices. *SIAM Journal on Matrix Analysis and Applications*, 31(3), 2010.
- [BP94] Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Number 9 in Classics in Applied Mathematics. Cambridge University Press, 1994.
- [BP98] Serguey Brin and Larry Page. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1-7):107–117, 1998. Proc. 17th International World Wide Web Conference.
- [BPMM03] Samuel Bernard, Laurent Pujo-Menjouet, and Michael C. Mackey. Analysis of cell kinetics using a cell division marker: Mathematical modeling of experimental data. *Biophys. J.*, 84:3414–3424, 2003.
- [BRR05] Allan Borodin, Gareth O. Roberts, and Jeffrey S. Rosenthal. Link analysis ranking: algorithms, theory, and experiments. *ACM Transactions on Internet Technology*, 5:231–297, 2005.

- [BS06] Mokhtar S. Bazaraa and Hanif D Sherali. *Nonlinear Programming: Theory and Algorithms*, 3rd ed. Jon Wiley, New York, 2006.
- [BS07] Holger Boche and Martin Schubert. Multiuser interference balancing for general interference functions - a convergence analysis. In *IEEE International Conference on Communications*, pages 4664–4669, 2007.
- [BT70] F. J. Burns and I. F. Tannock. On the existence of a G_0 phase in the cell cycle. *Cell Tissue Kinet.*, 3:321–334, 1970.
- [BT00] Vincent D. Blondel and John N. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000.
- [BTN01] Aharon Ben-Tal and Arkadi Nemirovski. *Lectures on Modern Convex Optimization, Analysis, Algorithms and Engineering Applications*. MPS/SIAM Series on Optimization. SIAM, 2001.
- [Bus73] Peter J. Bushell. Hilbert’s metric and positive contraction mappings in a banach space. *Archive for Rational Mechanics and Analysis*, 52:330–338, 1973.
- [BV04] Paolo Boldi and Sebastiano Vigna. The webgraph framework i: Compression techniques. In *Proc. of the 30th International World Wide Web Conference*, pages 595–601. ACM Press, 2004.
- [BYCL05] Ricardo A. Baeza-Yates, Carlos Castillo, and Vicente López. Pagerank increase under different collusion topologies. In *Proc. 1st International Workshop on Adversarial Information Retrieval on the Web*, pages 17–24, 2005.
- [BYRN99] Ricardo Baeza-Yates and Berthier Ribeiro-Neto. *Modern Information Retrieval*. Addison Wesley, 1999.
- [Cas06] Hal Caswell. *Matrix Population Models*. Encyclopedia of Environmetrics. John Wiley & Sons, Ltd, 2006.
- [CD10] Carlos Castillo and Brian D. Davison. Adversarial web search. In *Foundations and Trends in Information Retrieval*, volume 4, pages 377–486. Now Publishers, 2010.
- [CDW75] Jane Cullum, W. E. Donath, and P. Wolfe. The minimization of certain nondifferentiable sums of eigenvalues of symmetric matrices. In *Nondifferentiable Optimization*, volume 3 of *Mathematical Programming Studies*, pages 35–55. Springer Berlin Heidelberg, 1975.
- [CGL09] Jean Clairambault, Stéphane Gaubert, and Thomas Lepoutre. Comparison of perron and floquet eigenvalues in age structured cell division models. *Mathematical Modelling of Natural Phenomena*, 4:183–209, 2009.
- [CGL11] Jean Clairambault, Stéphane Gaubert, and Thomas Lepoutre. Circadian rhythm and cell population growth. *Mathematical and Computer Modelling*, 53:1558–1567, 2011.

- [CGP07] Jean Clairambault, Stéphane Gaubert, and Benoît Perthame. An inequality for the perron and floquet eigenvalues of monotone differential systems and age structured equations. *C. R. Acad. Sci. (Paris) Ser. I Mathématique*, 345:549–554, 2007.
- [CHS⁺04] David Cunningham, Yves Humblet, Salvatore Siena, David Khayat, Harry Bleiberg, Armando Santoro, Danny Bets, Matthias Mueser, Andreas Harstrick, Chris Verslype, Ian Chau, and Eric Van Cutsem. Cetuximab monotherapy and cetuximab plus irinotecan in irinotecan-refractory metastatic colorectal cancer. *N. Engl. J. Med.*, 351:337–345, 2004.
- [CJB09] Balázs Csanád Csáji, Raphaël M. Jungers, and Vincent D. Blondel. Pagerank optimization by edge selection. Eprint arXiv:0911.2280v1, 2009.
- [CJB10] Balázs Csanád Csáji, Raphaël M. Jungers, and Vincent D. Blondel. Pagerank optimization in polynomial time by stochastic shortest path reformulation. In *Proc. 21st International Conference on Algorithmic Learning Theory*, volume 6331 of *Lecture Notes in Computer Science*, pages 89–103. Springer, 2010.
- [CKM⁺09] A.A. Cohen, T. Kalisky, A. Mayo, N. Geva-Zatorsky, T. Danon, I. Issaeva, R.B. Kopito, N. Perzov, R. Milo, A. Sigal, and U. Alon. Protein dynamics in individual human cells: Experiment and theory. *PLoS one*, 4:1–12, 2009. e4901.
- [CL96] Steven Cox and Robert Lipton. Extremal eigenvalue problems for two-phase conductors. *Archive for Rational Mechanics and Analysis*, 136:101–117, 1996.
- [Cla07] Jean Clairambault. Modelling oxaliplatin drug delivery to circadian rhythm in drug metabolism and host tolerance. *Adv. Drug Deliv. Rev.*, 59:1054–1068, 2007.
- [Cla08] Jean Clairambault. A step toward optimization of cancer therapeutics. physiologically based modelling of circadian control on cell proliferation. *IEEE-EMB Magazine*, 27:20–24, 2008.
- [Cla09] Jean Clairambault. Modelling physiological and pharmacological control on cell proliferation to optimise cancer treatments. *Mathematical Modelling of Natural Phenomena*, 4:12–67, 2009.
- [Cla11] Jean Clairambault. Optimising cancer pharmacotherapeutics using mathematical modelling and a systems biology approach. *Personalized Medicine*, 8:271–286, 2011.
- [CLF09] D. Chaffey, C. Lake, and A. Friedlein. *Search Engine Optimization - Best Practice Guide*. Econsultancy.com Ltd, 2009.
- [CLMP03] Jean Clairambault, Béatrice Laroche, Stéphane Mischler, and Benoît Perthame. A mathematical model of the cell cycle and its control. Technical report, INRIA, Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France), 2003.
- [CLP11] Jean Clairambault, Alexander Lorz, and Benoît Perthame. Populational adaptive evolution and resistance to therapy. Submitted, 2011.

- [CMP06] Jean Clairambault, Philippe Michel, and Benoît Perthame. Circadian rhythm and tumour growth. *C. R. Acad. Sci. (Paris) Ser. I Mathématique (Équations aux dérivées partielles)*, 342:17–22, 2006.
- [CMP07] Jean Clairambault, Philippe Michel, and Benoît Perthame. A model of the cell cycle and its circadian control. In A. Deutsch, L. Brusch, H. Byrne, G. de Vries, and J. Herzog, editors, *Mathematical Modeling of Biological Systems, Volume I: Cellular Biophysics, Regulatory Networks, Development, Biomedicine, and Data Analysis*, pages 239–251. Birkhäuser, Boston, 2007.
- [CMTU01] G. Chiorino, J.A.J. Metz, D. Tomasoni, and P. Ubezio. Desynchronization rate in cell populations: mathematical modeling and experimental data. *J. Theor. Biol.*, 208:185–199, 2001.
- [Col42] Lothar Collatz. Einschliessungssatz für die charakteristischen zahlen von matrisen. *Mathematische Zeitschrift*, 48:221–226, 1942.
- [CT80] Michael G. Crandall and Luc Tartar. Some relations between nonexpansive and order preserving mappings. *Proceedings of the American Mathematical Society*, 78(3):385–390, 1980.
- [CZ91] Yair Censor and Stavros A. Zenios. Interval-constrained matrix balancing. *Linear Algebra and its Applications*, 150(0):393–421, 1991.
- [DDP08] P. Dua, V. Dua, and E.N. Pistikopoulos. Optimal delivery of chemotherapeutic agents in cancer. *Comp. Chem. Eng.*, 32:99–107, 2008.
- [DH89] R.B. Diasio and B.E. Harris. Clinical pharmacology of 5-fluorouracil. *Clin. Pharmacokinet.*, 16:215–237, 1989.
- [DHH⁺02] Chris Ding, Xiaofeng He, Parry Husbands, Hongyuan Zha, and Horst D. Simon. Pagerank, hits and a unified framework for link analysis. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 353–354, New York, NY, USA, 2002. ACM.
- [DK92] Alejandro R. Díaz and Noboru Kikuchi. Solutions to shape and topology eigenvalue optimization problems using a homogenization method. *International Journal for Numerical Methods in Engineering*, 35(7):1487–1502, 1992.
- [dKNvD08] Christobald de Kerchove, Laure Ninove, and Paul van Dooren. Maximizing pagerank via outlinks. *Linear Algebra and its Applications*, 429(5-6):1254–1276, 2008.
- [DL07] Jean-Charles Delvenne and Anne-Sophie Libert. Centrality measures and thermodynamic formalism for complex networks. *ArXiv e-print:0710.3972*, 2007.
- [DLL08] Michel De Lara and Doyen Luc. *Sustainable Management of Natural Resources - Mathematical Models and Methods*. Environmental Science and Engineering. Springer-Verlag Berlin Heidelberg, 2008.

- [dLMS09] Alberto d’Onofrio, Urszula Ledzewicz, Helmut Maurer, and Heinz Schättler. On optimal delivery of combination therapy for tumors. *Math. Biosci.*, 222:13–26, 2009.
- [DN84] Emeric Deutsch and Michael Neumann. Derivatives of the perron root at an essentially nonnegative matrix and the group inverse of an m-matrix. *Journal of Mathematical Analysis and Applications*, 102(1):1–29, 1984.
- [DN85] Emeric Deutsch and Michael Neumann. On the first and second order derivatives of the perron vector. *Linear algebra and its applications*, 71:57–76, 1985.
- [EHRS85] B. Curtis Eaves, Alan J. Hoffman, Uriel G. Rothblum, and Hans Schneider. Line-sum-symmetric scalings of square nonnegative matrices. In *Mathematical Programming Essays in Honor of George B. Dantzig Part II*, volume 25 of *Mathematical Programming Studies*, pages 124–141. Springer Berlin Heidelberg, 1985.
- [FABG13] Olivier Fercoq, Marianne Akian, Mustapha Bouhtou, and Stéphane Gaubert. Ergodic control and polyhedral approaches to Pagerank optimization. *IEEE Trans. Automatic Control*, 2013. arXiv preprint:1011.2348.
- [Fea10] John Fearnley. Exponential lower bounds for policy iteration. In *Automata, Languages and Programming*, volume 6199 of *Lecture Notes in Computer Science*, pages 551–562. Springer Berlin / Heidelberg, 2010.
- [Fer12a] Olivier Fercoq. Convergence of Tomlin’s HOTS algorithm. *arXiv preprint:1205.6727*, 2012.
- [Fer12b] Olivier Fercoq. Perron vector optimization applied to search engines. In *6th International conference on NETWORK Games, CONTROL and OPTimization (Net-goop)*, 2012. ArXiv preprint:1111.2234.
- [FG94] Shmuel Friedland and Leonid Gurvits. An upper bound for the real part of nonmaximal eigenvalues of nonnegative irreducible matrices. *SIAM Journal on Matrix Analysis and Applications*, 15(3):1015–1017, 1994.
- [FGH11] Shmuel Friedland, Stéphane Gaubert, and Lixing Han. Perron-frobenius theorem for nonnegative multilinear forms and extensions. *Linear Algebra and its Applications*, 2011. doi:10.1016/j.laa.2011.02.042.
- [Fil91] James Allen Fill. Eigenvalue bounds on convergence to stationarity for nonreversible markov chains, with an application to the exclusion process. *The Annals of Applied Probability*, 1(1):62–87, 1991.
- [FP00] K. Renee Fister and John Carl Panetta. Optimal control applied to cell-cycle-specific cancer chemotherapy. *SIAM Journal on Applied Mathematics*, 60(3):1059–1072, 2000.
- [Fri88] Shmuel Friedland. Bounds on the spectral radius of graphs with e edges. *Linear Algebra and its Applications*, 101:81–86, 1988.

- [Fri92] Shmuel Friedland. Lower bounds for the first eigenvalue of certain m-matrices associated with graphs. *Linear Algebra and its Applications*, 172:71–84, 1992.
- [Gat09] Robert A. Gatenby. A change of strategy in the war on cancer. *Nature*, 459:508–509, 2009.
- [GB82] Eli M Gafni and Dimitri P. Bertsekas. Convergence of a gradient projection method. Technical report, LIDS, MIT, 1982.
- [GCRG⁺08] A. Gréchez-Cassiau, B. Rayet, F. Guillaumond, M. Teboul, and F. Delaunay. The circadian clock component *bmal1* is a critical regulator of *p21(waf1/cip1)* expression and hepatocyte proliferation. *J. Biol. Chem.*, 283:4535–42, 2008.
- [GG04a] Stéphane Gaubert and Jeremy Gunawardena. Existence of the cycle time for some subtopical functions. Preprint, 2004.
- [GG04b] Stéphane Gaubert and Jeremy Gunawardena. The perron-frobenius theorem for homogeneous monotone functions. *Transactions of the AMS*, pages 4931–4950, 2004.
- [GGM05a] Zoltán Gyöngyi and Hector Garcia-Molina. Link spam alliances. In *VLDB '05: Proc. 31st international conference on Very large data bases*, pages 517–528. VLDB Endowment, 2005.
- [GGM05b] Zoltan Gyongyi and Hector Garcia-Molina. Web spam taxonomy. In *First International Workshop on Adversarial Information Retrieval on the Web (AIRWeb)*, 2005.
- [GGMP04] Zoltán Gyöngyi, Hector Garcia-Molina, and Jan Pedersen. Combating web spam with trustrank. In *Proc. 30th international conference on Very large data bases*, volume 30, pages 576–587. VLDB Endowment, 2004.
- [GGTW11] Pierre Gabriel, Shawn P. Garbett, Darren R. Tyson, and Glenn F. Webb. The contribution of age structure to cell population responses to targeted therapeutics. Submitted, 2011.
- [GL89] Gene H. Golub and Charles F. Van Loan. *Matrix computations*. Johns Hopkins University Press, Baltimore, Md, 1989.
- [GLD00] Robert Givan, Sonia Leach, and Thomas Dean. Bounded-parameter markov decision processes. *Artificial Intelligence*, 122(1-2):71–109, 2000.
- [GLM09] Anjela Y Govan, Amy N Langville, and Carl D Meyer. Offense-defense approach to ranking team sports. *Journal of Quantitative Analysis in Sports*, 5(1):1151, 2009.
- [GLS88] Martin Groetschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, 1988.
- [GOP02] Mats Gyllenberg, Andrei Osipov, and Lassi Päivärinta. The inverse problem of linear age-structured population dynamics. *J. Evol. Equations*, 2:223–239, 2002.

- [Gov08] Anjela Yuryevna Govan. *Ranking Theory with Application to Popular Sports*. PhD thesis, North Carolina State University, 2008.
- [GS07] Qingqing Gan and Torsten Suel. Improving web spam classifiers using link structure. In *Proc. 3rd inter. workshop on Adversarial information retrieval on the web AIRWeb '07*, pages 17–20. ACM (NY), 2007.
- [GSGF02] Robert A. Gatenby, Ariosto S. Silva, Robert J. Gillies, and B. Roy Frieden. Adaptive therapy. *Cancer Research*, 69:4894–4903, 2002.
- [HA89] Raphael T. Haftka and Howard M. Adelman. Recent developments in structural sensitivity analysis. *Structural and Multidisciplinary Optimization*, 1:137–151, 1989.
- [Har71] Darald J. Hartfiel. Concerning diagonal similarity of irreducible matrices. *Proceedings of the American Mathematical Society*, 30:413–425, 1971.
- [HDJ11] Romain Hollanders, Jean-Charles Delvenne, and Raphaël M. Jungers. Policy iteration is well suited to optimize pagerank. *arXiv preprint:1108.3779*, 2011.
- [HJ90] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, 1990. Corrected reprint of the 1985 original.
- [HNT99] Lixing Han, Michael Neumann, and Michael Tsatsomeros. Spectral radii of fixed frobenius norm perturbations of nonnegative matrices. *SIAM J. Matrix Anal. Appl.*, 21(1):79–92, 1999.
- [HWAW07] P. Hinow, S.E. Wang, C.L. Arteaga, and G.F. Webb. A mathematical model separates quantitatively the cytostatic and cytotoxic effects of a her2 tyrosine kinase inhibitor. *Theoretical Biology and Medical Modelling*, 4:14, 2007. doi:10.1186/1742-4682-4-14.
- [IT09] Hideaki Ishii and Roberto Tempo. Computing the pagerank variation for fragile web data. *SICE J. of Control, Measurement, and System Integration*, 2(1):1–9, 2009.
- [IT10] Hideaki Ishii and Roberto Tempo. Distributed randomized algorithms for the pagerank computation. *Automatic Control, IEEE Transactions on*, 55(9):1987–2002, 2010.
- [ITB12] Hideaki Ishii, Roberto Tempo, and Er-Wei Bai. A web aggregation approach for distributed randomized pagerank algorithms. *Automatic Control, IEEE Transactions on*, PP, 2012. In press.
- [JPS00] Charles R. Johnson, Joel Pitkin, and David P. Stanford. Line-sum symmetry via the DomEig algorithm. *Computational Optimization and Applications*, 17:5–10, 2000.
- [JSOvdD94] Charles R. Johnson, David P. Stanford, Dale Olesky, and Pauline van den Driessche. Dominant eigenvalues under trace-preserving diagonal perturbations. *Linear Algebra and its Applications*, 212-213(0):415–435, 1994.

- [Kat66] Tosio Kato. *Perturbation Theory for Linear Operators*. Springer-Verlag Berlin and Heidelberg GmbH & Co. K, 1966.
- [KB04] Michael B. Kastan and Jiri Bartek. Cell-cycle checkpoints and cancer. *Nature*, 432:316–323, 2004.
- [Kee93] James P. Keener. The perron-frobenius theorem and the ranking of football teams. *SIAM Review*, 35(1):80–93, 1993.
- [KHG06] Branislav Kveton, Milos Hauskrecht, and Carlos Guestrin. Solving factored mdps with hybrid state and action variables. *Journal of Artificial Intelligence Research*, 27:153–201, 2006.
- [Kin61] JFC Kingman. A convexity property of positive matrices. *The Quarterly Journal of Mathematics*, 12(1):283–284, 1961.
- [KKA06] Yuri Kheifetz, Yuri Kogan, and Zvia Agur. Long-range predictability in models of cell populations subjected to phase-specific drugs: growth-rate approximation using properties of positive compact operators. *Math. Models Methods Appl. Sci.*, 16:1155–1172, 2006.
- [Kle99] Jon Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46:604–632, 1999.
- [Kni08] Philip A. Knight. The sinkhorn-knopp algorithm: Convergence and applications. *SIAM J. on Matrix Analysis and Applications*, 30(1):261–275, 2008.
- [KR06] Vijay Krishnan and Rashmi Raj. Web spam detection with anti-trust rank. In *Proc. ACM SIGIR workshop on Adversarial Information Retrieval on the Web*, 2006.
- [KS67] Paul Knopp and Richard Sinkhorn. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific J. Math.*, 21(2):343–348, 1967.
- [KS06] Marek Kimmel and Andrzej Świerniak. Control theory approach to cancer chemotherapy: Benefiting from phase dependence and overcoming drug resistance. In Avner Friedman, editor, *Tutorials in Mathematical Biosciences III*, volume 1872 of *Lecture Notes in Mathematics*, pages 185–221. Springer Berlin / Heidelberg, 2006.
- [KYS⁺01] A. Kramer, F.C. Yang, P. Snodgrass, X. Li, and T.E. Scammell. Regulation of daily locomotor activity and sleep by hypothalamic EGF receptor signaling. *Science*, 294:2511–15, 2001.
- [L02] Francis Lévi. Cancer chronotherapeutics. *Special issue of Chronobiology International*, 19:1–19, 2002.
- [L06] Francis Lévi. Chronotherapeutics: the relevance of timing in cancer therapy. *Cancer Causes Control*, 17:611–621, 2006.
- [L08] Francis Lévi. The circadian timing system: A coordinator of life processes. implications for the rhythmic delivery of cancer therapeutics. *IEEE-EMB Magazine*, 27:17–20, 2008.

- [LACG08] Francis Lévi, Atilla Altinok, Jean Clairambault, and Albert Goldbeter. Implications of circadian clocks for the rhythmic delivery of cancer therapeutics. *Phil. Trans. Roy. Soc. A*, 366:3575–3598, 2008.
- [Lai64] Anna Kane Laird. Dynamics of tumour growth. *Br. J. Cancer*, 13:490–502, 1964.
- [LBK⁺07] H. Lodish, A. Berk, C.A. Kaiser, M. Krieger, M.P. Scotta, A. Bretscher, H. Ploegh, and P.T. Matsudaira. *Molecular Cell Biology*. W.H. Freeman, New York, 6th edition, 2007.
- [Lem01] Claude Lemaréchal. Lagrangian relaxation. In M. Jünger and D. Naddef, editors, *Computational Combinatorial Optimization*, pages 115–160. Springer Verlag, Heidelberg, 2001.
- [Lep09] Thomas Lepoutre. *Analyse et modélisation de phénomènes de croissance et mouvement issus de la biologie*. PhD thesis, Université Pierre-et-Marie-Curie (Paris VI), 2009.
- [LM00] Ronny Lempel and Shlomo Moran. The stochastic approach for link-structure analysis (SALSA) and the TKC effect. *Computer Networks*, 33(1-6):387–401, 2000.
- [LM05a] Amy N. Langville and Carl D. Meyer. A survey of eigenvector methods for web information retrieval. *SIAM Review*, 47(1):135–161, 2005.
- [LM05b] Amy N. Langville and Carl D. Meyer. A survey of eigenvector methods for web information retrieval. *SIAM Review*, 47(1):135–161, 2005.
- [LM06] Amy N. Langville and Carl D. Meyer. *Google’s PageRank and beyond: the science of search engine rankings*. Princeton University Press, 2006.
- [LMS11] Urszula Ledzewicz, Helmut Maurer, and Heinz Schaettler. Optimal and suboptimal protocols for a mathematical model for tumor anti-angiogenesis in combination with chemotherapy. *Mathematical Biosciences and Engineering*, 8(2):307–323, 2011.
- [LO96] Adrian S. Lewis and Michael L. Overton. Eigenvalue optimization. *Acta Numerica*, 5:149–190, 1996.
- [LOD⁺10] Francis Lévi, Alper Okyar, Sandrine Dulong, P.F. Innominato, and Jean Clairambault. Circadian timing in cancer treatments. *Annual Review of Pharmacology and Toxicology*, 50:377–421, 2010.
- [Log08] Dmitrii O. Logofet. Convexity in projection matrices: Projection to a calibration problem. *Ecological Modelling*, 216(2):217–228, 2008.
- [LS92] Ehud Lehrer and Sylvain Sorin. A uniform tauberian theorem in dynamic programming. *Mathematics of Operations Research*, 17(2):303–307, 1992.
- [LS07a] Urszula Ledzewicz and Heinz Schättler. Optimal controls for a model with pharmacokinetics maximizing bone marrow in cancer chemotherapy. *Math Biosci*, 206:320–342, 2007.

- [LS07b] Francis Lévi and Ueli Schibler. Circadian rhythms: Mechanisms and therapeutic implications. *Ann. Rev. Pharmacol. Toxicol.*, 47:493–528, 2007.
- [LT92] Zhi Quan Luo and Paul Tseng. On the convergence of the coordinate descent method for convex differentiable minimization. *Journal of Optimization Theory and Applications*, 72(1):7–35, 1992.
- [MA97] Harley H McAdams and Adam Arkin. Stochastic mechanisms in gene expression. *Proc. Natl. Acad. Sci. USA*, 31:814–819, 1997.
- [Mat96] Tomomi Matsui. NP-hardness of linear multiplicative programming and related problems. *Journal of Global Optimization*, 9:113–119, 1996.
- [Mat04] Fabien Mathieu. *Graphes du Web, Mesures d'importance à la PageRank*. PhD thesis, Université Montpellier II, 2004.
- [May94] Günter Mayer. Result verification for eigenvectors and eigenvalues. In *Proceedings of the IMACS-GAMM international workshop*, pages 209–276. Elsevier, 1994.
- [McK26] A.G. McKendrick. Applications of mathematics to medical problems. *Proc. Edinburgh Math. Soc.*, 54:98–130, 1926.
- [MD86] J.A.J. Metz and Odo Diekmann. *The dynamics of physiologically structured populations*, volume 68 of *Lecture notes in biomathematics*. Springer, New York, 1986.
- [Men76] Roy Mendelsohn. Optimization problems associated with a Leslie matrix. *The American Naturalist*, 110(973):339–349, 1976.
- [MF96] John M. McNamara and Par Forslund. Divorce rates in birds: Predictions from an optimization model. *The American Naturalist*, 147(4):609–640, 1996.
- [MMP05] Philippe Michel, Stéphane Mischler, and Benoît Perthame. The entropy structure of models of structured population dynamics. general relative entropy inequality: an illustration on growth models. *J. Math. Pures et Appl.*, 84:1235–1260, 2005.
- [Mor06] David Morgan. *The Cell Cycle: Principles of Control*. Primers in Biology series. Oxford University Press, 2006.
- [MRS⁺01] Joel C. Miller, Gregory Rae, Fred Schaefer, Lesley A. Ward, Thomas LoFaro, and Ayman Farahat. Modifications of kleinberg’s hits algorithm using matrix exponentiation and web log records. In *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2001.
- [MRS08] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schtze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.

- [MS88] Carl D. Meyer and G. W. Stewart. Derivatives and perturbations of eigenvectors. *SIAM J. on Numerical Analysis*, 25(3):679–691, 1988.
- [Mur90] John M. Murray. Some optimal control problems in cancer chemotherapy with a toxicity limit. *Math Biosci*, 100:49–67, 1990.
- [MV06] Fabien Mathieu and Laurent Viennot. Local aspects of the global ranking of web pages. In *6th International Workshop on Innovative Internet Community Systems (I2CS)*, pages 1–10, Neuchâtel, 2006.
- [MWB⁺00] M.-C. Mormont, J. Waterhouse, P. Bleuzen, S. Giacchetti, A. ami, A. Bogdan, J. Lellouch, J.-L. Misset, Y. Touitou, and F. Lévi. Marked 24-h rest-activity rhythms are associated with better quality of life, better response and longer survival in patients with metastatic colorectal cancer and good performance status. *Clin. Cancer Res.*, pages 3038–3045, 2000.
- [MYM⁺03] T. Matsuo, S. Yamaguchi, S. Mitsui, A. Emi, F. Shimoda, and H. Okamura. Control mechanism of the circadian clock for timing of cell division in vivo. *Science*, 302:255–259, 2003.
- [NA08] Danil Nemirovsky and Konstantin Avrachenkov. Weighted pagerank: Cluster-related weights. In Ellen M. Voorhees and Lori P. Buckland, editors, *Proc. 17th Text REtrieval Conference, TREC’08*. National Institute of Standards and Technology (NIST), 2008.
- [Nel76] R B Nelson. Simplified calculation of eigenvector derivatives. *AIAA Journal*, 14:1201–1205, 1976.
- [Nin08] Laure Ninove. *Dominant Vectors of Nonnegative Matrices*. Phd thesis, Université Catholique de Louvain, 2008.
- [NNMF06] Alexandros Ntoulas, Marc Najork, Mark Manasse, and Dennis Fetterly. Detecting spam web pages through content analysis. In *Proc. 15th international conference on World Wide Web*, pages 83–92, New York, 2006. ACM.
- [NP09] Aleksandr Nazin and Boris Polyak. Adaptive randomized algorithm for finding eigenvector of stochastic matrix with application to pagerank. In *Proc. 48th IEEE Conference on Decision and Control (CDC/CCC)*, pages 127–132, 2009.
- [NP11] Yurii Nesterov and Vladimir Yu. Protasov. Optimizing the spectral radius. preprint, 2011.
- [Nus88] Roger D. Nussbaum. Hilbert’s projective metric and iterated nonlinear maps. In *Memoirs of the American Mathematical Society*, volume 75 (391). AMS, 1988.
- [NW99] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Nueva York, EUA : Springer, 1999.
- [Ost55] Alexander Ostrowski. über normen von matrizen. *Mathematische Zeitschrift*, 63:2–18, 1955.

- [Ove91] Michael L. Overton. Large-scale optimization of eigenvalues. *SIAM J. on Optimization*, 2(1):88–120, 1991.
- [OW88] Michael L. Overton and Robert S. Womersley. On minimizing the spectral radius of a nonsymmetric matrix function: optimality conditions and duality theory. *SIAM J. on Matrix Analysis and Applications*, 9:473–498, 1988.
- [PA95] John Carl Panetta and John A. Adam. A mathematical model of cell-specific chemotherapy. *Math Comput Modelling*, 22:67, 1995.
- [Pan97] John Carl Panetta. A mathematical model of breast and ovarian cancer treated with paclitaxel. *Math. Biosci.*, 146(2):89–113, 1997.
- [Per07] Benoît Perthame. *Transport Equations in Biology*. Frontiers in Mathematics series. Birkhäuser, Boston, 2007.
- [PLK⁺93] G.J. Peters, J. Lankelma, R.M. Kok, P. Noordhuis, C.J. van Groenigen, and C.L. van der Wilt et al. Prolonged retention of high concentrations of 5-fluorouracil in human and murine tumors as compared with plasma. *Cancer Chemother. Pharmacol.*, 31:269–276, 1993.
- [Pol97] Elijah Polak. *Optimization: algorithms and consistent approximations*. Springer-Verlag New York, 1997.
- [PP02] Olivier Pironneau and Elijah Polak. Consistent approximations and approximate functions and gradients in optimal control. *SIAM J. Control Optim.*, 41:487–510, 2002.
- [PPdS95] F.L. Pereira, C.E. Pedreira, and J.B. de Sousa. A new optimization based approach to experimental combination chemotherapy. *Frontiers Med. Biol. Engng.*, 6(4):257–268, 1995.
- [PPJ73] Beresford N. Parlett and William G. Poole Jr. A geometric theory for the QR, LU and power iterations. *SIAM J. on Numerical Analysis*, 10(2):389–412, 1973.
- [Pro06] Academic Web Link Database Project. New Zealand university web sites. <http://cybermetrics.wlv.ac.uk/database/index.html>, January 2006.
- [PT87] Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987.
- [Puk06] Friedrich Pukelsheim. *Optimal Design of Experiments*, volume 50 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2006.
- [Put94] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.
- [RW82] Uriel G. Rothblum and Peter Whittle. Growth optimality for branching markov decision chains. *Mathematics of Operations Research*, 7(4):582–601, 1982.
- [RW97] R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*, volume 317 of *Grundlehren der math. Wiss.* Springer, 1997.

- [RW02] Steven M. Reppert and David R Weaver. Coordination of circadian timing in mammals. *Nature*, 418:935–941, 2002.
- [Saa87] Thomas L. Saaty. Rank according to perron: A new insight. *Mathematics Magazine*, 60(4):211–213, 1987.
- [Sch89] Michael H. Schneider. Matrix scaling, entropy minimization, and conjugate duality (i): existence conditions. *Linear Algebra and its Applications*, 114-115:785–813, 1989. Special Issue Dedicated to Alan J. Hoffman.
- [Sch90] Michael H. Schneider. Matrix scaling, entropy minimization and conjugate duality (ii): The dual problem. *Math. Prog.*, 48:103–124, 1990.
- [SF95] Alexander Shapiro and Michael K. H. Fan. On eigenvalue optimization. *SIAM J. on Optimization*, 5(3):552–569, 1995.
- [SM83] Gerard Salton and Michael J. McGill. *Introduction to modern information retrieval*. McGraw-Hill (New York), 1983.
- [Smi05] Warren D. Smith. Sinkhorn ratings, and new strongly polynomial time algorithms for sinkhorn balancing, perron eigenvectors, and markov chains. Available: <http://rangevoting.org/WarrenSmithPages/homepage/works.html>, 2005.
- [SS08] Vahid Shahrezaei and Peter G. Swain. Analytical distributions for stochastic gene expression. *Proc. Natl. Acad. Sci. USA*, 105:17256–17261, 2008.
- [SSC02] Ueli Schibler and Paolo Sassone-Corsi. A web of circadian pacemakers. *Cell*, 111:919–922, 2002.
- [SSKM⁺08] A. Sakaue-Sawano, H. Kurokawa, T. Morimura, A. Hanyu, H. Hama, H. Osawa, S. Kashiwagi, K. Fukami, T. Miyata, H. Miyoshi, T. Imamura, M. Ogawa, H. Masai, and A. Miyawaki. Visualizing spatiotemporal dynamics of multicellular cell-cycle progression. *Cell*, 32:487–498, 2008.
- [SSOH⁺08] A. Sakaue-Sawano, K. Ohtawa, H. Hama, M. Kawano, M. Ogawa, and A. Miyawaki. Tracing the silhouette of individual cells in S/G2/M phases with fluorescence. *Chem Biol.*, 15:1243–48, 2008.
- [STH⁺08] E. Sherer, E. Tocce, R.E. Hannemann, A.E. Rundell, and D. Ramkrishna. Identification of age-structured models: cell cycle phase transitions. *Biotechnol. Bioeng.*, 99:960–974, 2008.
- [STKA07] Hiroo Saito, Masashi Toyoda, Masaru Kitsuregawa, and Kazuyuki Aihara. A large-scale study of link spam detection by graph algorithms. In *Proc. 3rd inter. workshop on Adversarial information retrieval on the web AIRWeb '07*, pages 45–48. ACM New York, 2007.
- [SZ90] Michael H. Schneider and Stavros A. Zenios. A comparative study of algorithms for matrix balancing. *Operations Research*, 38(3):439–455, 1990.

- [TB97] John N. Tsitsiklis and Vincent D. Blondel. The lyapunov exponent and joint spectral radius of pairs of matrices are hard - when not impossible - to compute and to approximate. *Mathematics of Control, Signals, and Systems (MCSS)*, 10:31–40, 1997.
- [TFL11] Chee Wei Tan, Shmuel Friedland, and Steven H. Low. Spectrum management in multiuser cognitive wireless networks: Optimality and algorithm. *Selected Areas in Communications, IEEE Journal on*, 29(2):421–430, 2011.
- [Tom03] John A. Tomlin. A new paradigm for ranking pages on the world wide web. In *Proc. 12th inter. conf. on World Wide Web, WWW '03*, pages 350–355, New York, 2003. ACM.
- [Tse90] Paul Tseng. Solving h-horizon, stationary markov decision problems in time proportional to $\log(h)$. *Operations Research Letters*, 9(5):287–297, 1990.
- [Tut01] William Thomas Tutte. *Graph theory*, volume 21 of *Enc. of Mathematics and its Applications*. Cambridge University Press, 2001.
- [Vig09a] Guillaume Vigerel. *Propriétés asymptotiques des jeux répétés à somme nulle*. PhD thesis, Université Pierre et Marie Curie - Paris VI, 2009.
- [Vig09b] Sebastiano Vigna. Spectral ranking. *arXiv preprint:0912.0238*, 2009.
- [Vla12] Nikos Vlassis. Np-hardness of polytope m-matrix testing and related problems. *arXiv preprint:1206.2059*, 2012.
- [VLL00] Bert Vogelstein, David Lane, and Arnold J. Levine. Surfing the p53 network. *Nature*, 408:307–310, 2000.
- [VOA10] Minaya Villasana, Gabriela Ochoa, and Soraya Aguilar. Modeling and optimization of combined cytostatic and cytotoxic cancer chemotherapy. *Artificial Intelligence in Medicine*, 50(3):163–173, 2010.
- [VVB03] K. Vermeulen, D. R. Van Bockstaele, and Z. N. Berneman. The cell cycle: a review of regulation, deregulation and therapeutic targets in cancer. *Cell Prolif.*, 36:131–149, 2003.
- [web] The stanford webbase project. <http://diglib.stanford.edu:8091/~testbed/doc2/WebBase/>.
- [Web90] Glenn F. Webb. Resonance phenomena in cell population chemotherapy models. *Rocky Mountain J. Math*, 20(4):1195–1216, 1990.
- [Web92] Glenn F. Webb. A cell population model of periodic chemotherapy treatment. *Biomedical Modeling and Simulation*, pages 83–92, 1992.
- [web07] Webspam-uk2007, 2007. Yahoo! Research: "Web Spam Collections". Crawled by the Laboratory of Web Algorithmics, University of Milan.
- [WV07] Hans-Fabian Winkler and Silviu Vlasceanu. Quelle est la valeur d'un hyperlien ? Rapport d'enseignement d'approfondissement en recherche opérationnelle, Ecole Polytechnique, 2007.

-
- [XM10] Huan Xu and Shie Mannor. Distributionally robust markov decision processes. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2505–2513. Informs, 2010.
- [Ye05] Yinyu Ye. A new complexity result on solving the markov decision problem. *Mathematics of Operations Research*, 30(3):733–749, 2005.
- [Ye11] Yinyu Ye. The simplex and policy-iteration methods are strongly polynomial for the markov decision problem with a fixed discount rate. *Mathematics of Operations Research*, 36(4):593–603, 2011.
- [ZC91] Stavros A. Zenios and Yair Censor. Massively parallel row-action algorithms for some nonlinear transportation problems. *SIAM J. Optim*, 1:373–400, 1991.
- [Zie01] Gunter M. Ziegler. *Lectures on Polytopes (Graduate Texts in Mathematics)*. Springer, 2001.