

Ergodic Control and Polyhedral approaches to PageRank Optimization

Olivier Fercoq ^{*,‡} Marianne Akian ^{*} Mustapha Bouhtou [†]
Stéphane Gaubert ^{*}

September 20, 2011

Abstract

We study a general class of PageRank optimization problems which consist in finding an optimal outlink strategy for a web site subject to design constraints. We consider both a continuous problem, in which one can choose the intensity of a link, and a discrete one, in which in each page, there are obligatory links, facultative links and forbidden links. We show that the continuous problem, as well as its discrete variant when there are no constraints coupling different pages, can both be modeled by constrained Markov decision processes with ergodic reward, in which the webmaster determines the transition probabilities of websurfers. Although the number of actions turns out to be exponential, we show that an associated polytope of transition measures has a concise representation, from which we deduce that the continuous problem is solvable in polynomial time, and that the same is true for the discrete problem when there are no coupling constraints. We also provide efficient algorithms, adapted to very large networks. Then, we investigate the qualitative features of optimal outlink strategies, and identify in particular assumptions under which there exists a “master” page to which all controlled pages should point. We report numerical results on fragments of the real web graph.

^{*}INRIA Saclay and CMAP Ecole Polytechnique
`olivier.fercoq@inria.fr`
`marianne.akian@inria.fr`
`stephane.gaubert@inria.fr`

[†]France Télécom R & D
`mustapha.bouhtou@orange-ftgroup.com`

[‡]The doctoral work of the first author is supported by Orange Labs through the research contract CRE 3795 with INRIA

1 Introduction

The PageRank introduced by Brin and Page [1] 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 [2] for a discussion of the PageRank optimization methods 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 [3], [4].

The PageRank has motivated a number of works, dealing in particular with computational issues. Classically, the PageRank vector is computed by the power algorithm [1]. There has been a considerable work on designing new, more efficient approaches for its computation [5, 6]: Gauss-Seidel method [7], aggregation/disaggregation [6] or distributed randomized algorithms [8, 9]. Other active fields are the development of new ranking algorithms [10] or the study of the web graph [11].

The optimization of PageRank has been studied by several authors. Avrachenkov and Litvak analyzed in [12] the case of a single controlled page and determined an optimal strategy. In [13], 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 [14] 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 [15], Ninove developed a heuristic based on these theoretical results, which was experimentally shown to be efficient. In [16], 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 [17], (see also [18] 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.

In the present paper, we study 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 [19], its optimization does not seem to have 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 [16, 17], 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). When specialized to the discrete problem, this extends Theorem 1 of [17], which only applies to the case in which the utility function represents the PageRank of a single page. The proof of Theorem 4 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 1). This leads us to 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 [20]). We prove in Theorem 3, as an application of the theory of Khachiyan’s ellipsoid method (see [21]), that the ergodic control problem for well-described Markov decision process is polynomial time solvable (even in the multi-chain framework). Then, Theorem 4 follows as a direct corollary. We note that maximization or separation oracles have been previously considered in dynamic programming for different purposes (dealing with unknown parameters [22, 23], or approximating large scale problems [24]).

Proposition 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 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 [17], 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 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 13 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 [17] 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 7.3). On test instances, approximate optimality certificates show that the solution found by the heuristic is at most at 1.7% of the optimum.

Using the concept of mean reward before teleportation, we identify in Theorem 5 (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 [14], but in [14], 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 6).

We report numerical results on the web site of one of the authors (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 [25].

The paper is organized as follows. In Section 2, we introduce the general PageRank optimization problem. In Section 3, we give a concise description of the polytope of uniform transition probabilities. In Section 4, we show that every Markov decision process which admits such a concise description is polynomial time solvable (Theorem 3), and we deduce as a corollary our first main result, Theorem 4. Section 4.3 describes an efficient fixed point scheme for the resolution of the PageRank optimization problem with local constraints. In Section 5, we give the “master page” Theorem (Theorem 5). We deal with coupling constraints in Section 6. We give experimental results

on real data in Section 7.

2 PageRank optimization problems

2.1 Google's PageRank

We first recall the basic elements of the Google PageRank computation, see [1] and [6] for more information. 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

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

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

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

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 (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 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 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:

$$\max_{\pi, P} \{U(\pi, P) ; \pi = \pi P, \quad \pi \in \Sigma_n, \quad P \in \mathcal{P}\} \quad (2)$$

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 $P_{i\cdot}$ the

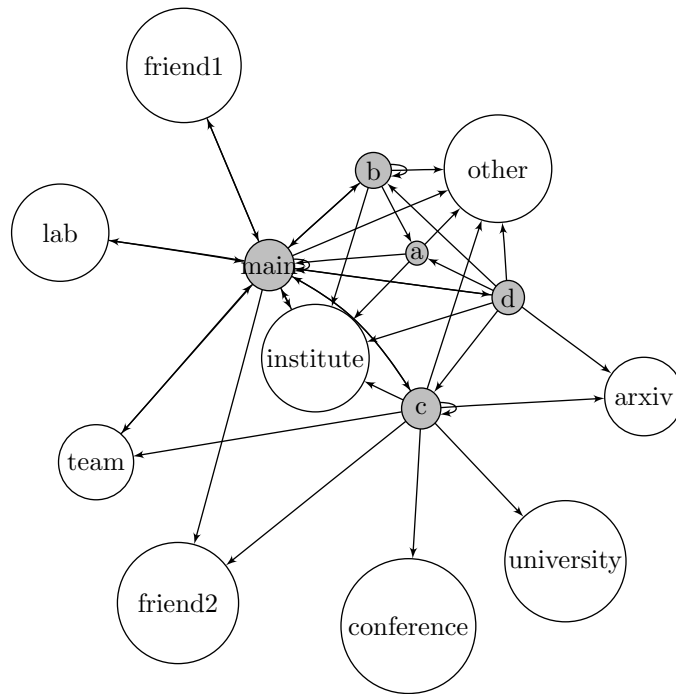


Figure 1: The web site of one of the authors (colored) and the surrounding sites (white). This 1500-page fragment of the web is aggregated for presentation, using the technique described in [6]. The sizes of the circles follow the log of their PageRank.

i th row of a matrix P . We shall distinguish *local constraints*, which can be expressed as $P_{i,\cdot} \in \mathcal{P}_i$, where $\mathcal{P}_i \subset \Sigma_n$ is a given subset, and *global constraints*, which couple several vectors $P_{i,\cdot}$. Thus, local constraints only involve the outlinks from a single page, whereas global constraints involve the outlinks from different pages. We shall consider the situation in which each \mathcal{P}_i is a polytope (or more generally an effective convex set).

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 [16] and [17], 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 (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:

$$\max_{\pi, P} \{U(\pi, P) \ ; \ \pi = \pi P, \ \pi \in \Sigma_n, \ P = P(J_1, \dots, J_n), \ J_i \subseteq \mathcal{F}_i, \ i \in [n]\} \quad (3)$$

Remark 1. Problem (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:

$$U(\pi, P) = \sum_{i \in [n]} \pi_i \sum_{j \in [n]} P_{i,j} r_{i,j} . \quad (4)$$

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

Remark 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). 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 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 $\sum_{j \in [n]} r'_{i,j} z_j = \sum_{j \in [n]} \sum_{l \in [n]} r'_{i,l} z_l P_{i,j}$ 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

$$\rho_{i,j} \geq 0, \forall i, j \in [n], \quad \sum_{i,j \in [n]} \rho_{i,j} = 1, \quad \sum_{k \in [n]} \rho_{k,i} = \sum_{j \in [n]} \rho_{i,j}, \forall i \in [n]. \quad (5)$$

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:

$$P = f(\rho), \quad P_{i,j} = \frac{\rho_{i,j}}{\sum_k \rho_{i,k}}, \quad \forall i, j \in [n]. \quad (6)$$

Proposition 1. *The function f defined by (6) sets up a birational transformation between the set of irreducible occupation measures (irreducible matrices satisfying (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 [26]). 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: $U(\pi, P) = \sum_{i,j \in [n]} \rho_{i,j} r_{i,j}$.

2.3 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 (2).

Proposition 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 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 $H_i^{(l)}$:

$$\sum_{j \in [n]} a_{i,j}^{(l)} P_{i,j} \leq b_i^{(l)}, \quad \forall i \in [n], \forall l \in L \quad (7)$$

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 1, and rewriting Inequalities (7) in the form

$$\sum_{j \in [n]} a_{i,j}^{(l)} \rho_{i,j} \leq b_i^{(l)} \sum_{k \in [n]} \rho_{i,k}, \quad \forall i \in [n], \forall l \in L \quad (8)$$

we see that ρ satisfies a family of constraints of the form (8), together with the inequalities (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 .

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 [14], setting mutual exclusion between two hyperlinks [17] or limiting the number of hyperlinks [17]. 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 7.3).

3 The polytope of uniform transition measures

In this section, we show that the polytope of uniform transition measures admits a concise representation (Theorem 1). The vertices of this polytope represent the action space of the Discrete PageRank Optimization problem (3). Theorem 1 is a key ingredient of the proof of the polynomial time character of this problem which will be given in the next section.

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 $\alpha = 1$. To get back to $\alpha < 1$, we use the relation $P_{i,j} = \alpha S_{i,j} + (1 - \alpha)z_j$ (see Remark 5 at the end of this section).

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,

$$\mathcal{D}_i = \{q \in \Sigma^n \mid \mathcal{O}_i \subseteq \text{supp}(q) \subseteq \mathcal{O}_i \cup \mathcal{F}_i, \\ q \text{ uniform probability measure on its support}\} \quad (9)$$

or if $\mathcal{O}_i = \emptyset$, it is possible to have no link at all and then to teleport with probability vector Z :

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

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 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:*

$$\forall j \in \mathcal{I}_i, \quad x_j = 0 \quad \forall j \in \mathcal{F}_i, x_j \leq x_{j_0} \quad (10a)$$

$$\forall j \in \mathcal{O}_i \setminus \{j_0\}, x_j = x_{j_0} \quad \forall j \in \mathcal{F}_i, x_j \geq 0 \quad (10b)$$

$$\sum_{j \in [n]} x_j = 1 \quad (10c)$$

Proof. Let \mathcal{S}_i be the polytope defined by Inequalities (10).

$(\mathcal{D}_i \subseteq \mathcal{S}_i)$: 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 (10b) and (10a) 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 (10b) and (10a) cannot be saturated simultaneously (see the previous paragraph), so the only way to saturate $|\mathcal{F}_i|$ inequalities is to saturate one of (10b) or (10a) 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 ([27], 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 $\{X \in \mathbb{R}^n | X_{j_0} = 1\}$. We can then construct the homogenization of H , $\text{homog}(H)$, which is the pointed cone with base H (see [27] for more details). Finally \mathcal{S}_i is the cross-section of $\text{homog}(H)$ with the hyperplane $\{x \in \mathbb{R}^n | \sum_{j \in [n]} x_j = 1\}$. \square

The result of the theorem implies in particular that $\text{co}(\mathcal{D}_i)$ is combinatorially equivalent to a hypercube, ie. that their face lattices are isomorphic [27].




Figure 2: Projection of the polytope of uniform transition measures with one obligatory link ($|\mathcal{O}_i| = 1$) and three facultative links ($|\mathcal{F}_i| = 3$).

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 .

Proposition 3. *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:

$$\sum_{j \in [n]} x_j = 1, \quad x_k \geq 0 \quad (11a)$$

$$\forall j \in \mathcal{I}_i \setminus \{k\}, \quad x_j = \frac{Z_j}{Z_k} x_k, \quad \forall j \in \mathcal{F}_i, \quad x_j \geq \frac{Z_j}{Z_k} x_k \quad (11b)$$

Proof. The proof follows the same sequence of arguments as the proof of Theorem 1. We just need to adapt it to Inequalities (11). \square

Proposition 4. *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. When there is no obligatory link, most of the admissible discrete transition probabilities are not extreme points of the polytope.

Remark 5. If we want to work with \mathcal{P}_i , the polytope of transition probabilities with damping factor α , we only need the relation $\mathcal{P}_i = \alpha \mathcal{S}_i + (1 - \alpha)z$ to get the actual inequalities. For instance, $x_j = x_{j_0}$ remains but $x_j \geq 0$ becomes $x_j \geq (1 - \alpha)z_j$.

4 Solving the PageRank Optimization Problem with local constraints

4.1 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.

A finite *Markov decision process* is a 4-uple $(I, (A_i)_{i \in I}, p, r)$ 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 ; $p : I \times \cup_{i \in I} (\{i\} \times A_i) \rightarrow \mathbb{R}_+$ 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 $t \geq 0$, $\nu_t(h_t) = \bar{\nu}(X_t)$.

Given an initial distribution μ representing the law of X_0 , the average cost infinite horizon Markov decision problem, also called *ergodic control problem*, consists in maximizing

$$\liminf_{T \rightarrow +\infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} r(X_t, \nu_t) \right) \quad (12)$$

where the maximum is taken over the set of randomized control strategies ν . Indeed, the supremum is the same if it is taken only over the set of randomized (or even deterministic) stationary feedback strategies (Theorem 9.1.8 in [28] 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, its value 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 (12) when starting from state i .

Proposition 5. *If there are only local constraints, ie. $\mathcal{P} = \prod_{i \in [n]} \mathcal{P}_i$, if for all $i \in [n]$, \mathcal{P}_i is a polytope and if the utility function is an income proportional to the frequency of clicks (4), then the continuous PageRank Optimization problem (2) 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 6. *The following dynamic programming equation:*

$$w_i + \psi = \max_{\nu \in \mathcal{P}_i} \nu(r_{i, \cdot} + w) \quad , \quad \forall i \in [n] \quad (13)$$

has a solution $w \in \mathbb{R}^n$ and $\psi \in \mathbb{R}$. The constant ψ is unique and is the value of the PageRank Optimization problem (2). An optimal strategy is obtained by selecting for each state i a maximizing $\nu \in \mathcal{P}_i$ in (13). The function w is often called the bias or the potential.

Proof. Theorem 8.4.3 in [28] applied to the unichain ergodic control problem of Proposition 5 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

Theorem 2. *The discrete PageRank Optimization problem (3) is equivalent to a continuous PageRank Optimization problem (2) in which the action set \mathcal{P}_i is defined by one of the polytopes described in Theorem 1 or Proposition 3 or 4, depending on the presence of obligatory links.*

Proof. Arguing as in the proof of Proposition 5, we get that the discrete PageRank Optimization problem (3) is equivalent to an ergodic control problem with state space $[n]$, in which the action set in state i is the discrete set \mathcal{D}_i defined in (9), and the rewards and transition probabilities are as in this proposition. The optimal solutions of the discrete PageRank Optimization problem coincide with the optimal stationary deterministic strategies. The analog of Equation (13) is now

$$w_i + \psi = \max_{\nu \in \text{co}(\mathcal{D}_i)} \nu(r_{i, \cdot} + w) \quad (14)$$

where $\text{co}(\mathcal{D}_i)$ is the convex hull of the set \mathcal{D}_i , i.e the polytope described in either Theorem 1 or Proposition 3 or 4. 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 \mathcal{D}_i or

$\text{co}(\mathcal{D}_i)$ have the same value. Moreover, an optimal strategy of the problem with the latter set of actions can be found by solving (14) and selecting a maximizing action ν in (14). Such an action may always be chosen in the set of extreme points of $\text{co}(\mathcal{D}_i)$ and these extreme points belong to \mathcal{D}_i (beware however that some points of \mathcal{D}_i may be not extreme). \square

4.2 Polynomial time solvability of well-described Markov decision problems

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 [20] 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 1).

However, we next establish a general result, Theorem 3 below, showing that the polynomial time solvability of ergodic control problems subsists when the action sets are implicitly defined. This is based on the combinatorial developments of the theory of Khachiyan’s ellipsoid method, by Groetschel, Lovász and Schrijver [21]. We refer the reader to the latter monograph for more background on the notions of strong separation oracles and well described polyhedra.

Definition 1 (Def. 6.2.2 of [21]). 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 $\mathcal{B} \in \mathbb{R}^n$ is a polyhedron with facet-complexity at most ϕ . The *encoding length* of \mathcal{B} is by definition $n + \phi$.

Definition 2 (Problem (2.1.4) of [21]). 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 1, we introduce the following notion.

Definition 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 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 Theorem 1: 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 (10) is not satisfied.

Theorem 3. *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. Consider the polyhedron \mathcal{G} consisting of the couples of vectors $(v, g) \in \mathbb{R}^I \times \mathbb{R}^I$ satisfying the constraints

$$\begin{aligned} g_i &\geq \sum_{j \in I} \sum_{l \in [L_i]} a_l Q_{i,j}^l g_j, \quad \forall i \in I, a \in A_i \\ v_i + g_i &\geq \sum_{l \in [L_i]} a_l R_i^l + \sum_{j \in I} \sum_{l \in [L_i]} a_l Q_{i,j}^l v_j, \quad \forall i \in I, a \in A_i. \end{aligned} \tag{15}$$

Theorem 9.3.8 in [28] 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 [21], 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 [21] 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 (15). 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 (15), 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 [21], 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 6. This argument also shows that the discounted problem is polynomial time solvable.

As a consequence of Theorems 2 and 3, we get

Theorem 4. *If there are only local constraints, if the utility function is a rational total income utility (4) and if the teleportation vector and damping factor are rational, then the discrete problem (3) can be solved in polynomial time and the continuous problem (2) with well-described action sets (Definition 1) can also be solved in polynomial time.*

Proof. Thanks to Theorem 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 3). 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 polytope as it is the case here (Lemma 6.5.1 in [21]).

By Proposition 5, the ergodic control problem associated to a continuous PageRank Optimization problem with well-described action sets satisfies the conditions of Theorem 3 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

Theorem 3 is mostly of theoretical interest, since its proof is based on the ellipsoid algorithm, which is slow. We however give in Section 4.3 a fast scalable algorithm for the present problem.

Example 1. Consider again the graph from Figure 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 2.3): each page can change up to 20 % of the initial transition probabilities. The result is represented in Figure 3.

Example 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 by squares. Facultative links are all other possible links from controlled pages.

4.3 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 Remark 14 in Section 5) 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 7. *Let T be the dynamic programming operator $\mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by*

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

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 (13).

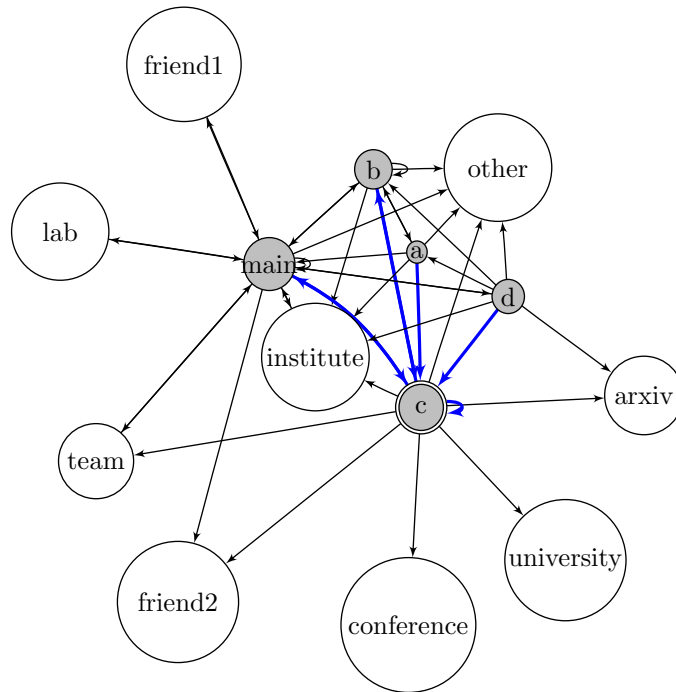


Figure 3: Web graph of Figure 1 optimized under local skeleton constraints. The optimal strategy consists in linking as much as possible to page "c" (actually, the page of a lecture), up to saturating the skeleton constraint. This page gains then a PageRank comparable to the one of the main page. The sum of the PageRank scores has been increased by 22.6%.




Figure 4: 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 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 5).

Proof. The set $\{\nu \text{ st } \alpha\nu + (1 - \alpha)z \in \mathcal{P}_i\}$ is a set of probability measures so it is clear that T is α -contracting. Let w be its fixed point. For all $i \in [n]$,

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

We get equation (13) with constant $(1 - \alpha)zw$. \square

Remark 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 3).

Remark 8. The fixed point found is just the mean reward before teleportation at the optimum (see Definition 4, Section 5) .

We can then solve the dynamic programming equation (13) and so the PageRank Optimization Problem (2) or (3) with local constraints by value iteration.

The algorithm starts with an initial potential function w , scans repeatedly the pages and updates w_i when i is the current page according to $w_i \leftarrow T_i(w)$ until convergence is reached. Then $(w, (1 - \alpha)zw)$ is solution of the ergodic dynamic programming equation (13) and 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.

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 3 and 4), so we only describe it in the other case. In Algorithm 1, 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.

Proposition 8. *When the constraints of the Discrete PageRank Optimization problem (3) are defined by obligatory, facultative and forbidden links, the greedy algorithm (Algorithm 1) started at page i returns $T_i(w)$ as defined in Proposition 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

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

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 9. A straightforward modification of the greedy algorithm can handle a upper or a lower limit on the number of links on a given page.

Proposition 9. *An ϵ approximation of the Discrete PageRank Optimization Problem (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 result should be compared to PageRank computation's complexity by the power method [5], which is $O\left(\frac{\log(\epsilon)}{\log(\alpha)} \sum_{i \in [n]} |\mathcal{O}_i| + |\mathcal{F}_i|\right)$.

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 [14], we shall use the mean reward before teleportation to study the optimal outlink strategies.

Definition 4. 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 10. *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). Then P is an optimal link strategy of the continuous PageRank Optimization problem (2) 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 $\forall i \in [n], v_i(P) + \pi(P)r' = \max_{\nu \in \mathcal{P}_i} \nu(v(P) + r'_i e)$. By Proposition 6, this means that $v(P)$ is the bias of Equation (13) and that P is an optimal outlink strategy. \square

Remark 10. Proposition 10 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 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 5 (Master Page). *Consider the Discrete PageRank Optimization problem (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 v be the mean reward before teleportation (Definition 4) at the optimum. Then any optimal link strategy must choose for every controlled page i all the facultative links (i, j) such that $v_j > \frac{v_i - r'_i}{\alpha}$*




Figure 5: Maximization of the sum of the PageRank values of the colored pages. Top: obligatory links; self links are forbidden; all other links 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.

and any combination of facultative links such that $v_j = \frac{v_i - r'_i}{\alpha}$. Moreover, all optimal link strategies are obtained in this way.

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. By Remark 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 1). Moreover, by Proposition 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 $\frac{1}{|J| + |\mathcal{O}_i|} \sum_{l \in J \cup \mathcal{O}_i} v_l + r'_i < v_j + r'_i$, 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. \square

The theorem is illustrated in Example 2 (Section 4) and Figure 5.

Example 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 12. If every controlled page is allowed to point to every page, as in Figures 3 and 4, 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 10 and Theorem 5 can be related to [3], 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 [4], 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 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.

Remark 14. 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 ([29] for instance), these are still incomplete crawlings when compared to the search engine's.

We denote by v and \tilde{v} the mean reward before teleportation of respectively the search engine's web graph and the truncated web graph. 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. We can easily show that if there are no page without outlink, then for all i in I , $|v_i - \tilde{v}_i| \leq \alpha^{R+1} \frac{2}{1-\alpha} \|\bar{r}\|_\infty$.

When there are pages without outlink, the problem is more technical. A possible approach to deal with it is to use the non-compensated PageRank [30].

6 PageRank Optimization with coupling constraints

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

From now on, we have studied discrete or continuous PageRank Optimization problems but only with local constraints. We consider in this section the following PageRank Optimization problem (2) with ergodic (linear in the occupation measure) coupling constraints:

$$\begin{aligned} \max_{\pi, P} \sum_{i,j} \pi_i P_{i,j} r_{i,j} \text{ st:} \\ \pi P = \pi, \pi \in \Sigma_n, P_{i,\cdot} \in \mathcal{P}_i, \forall i \in [n] \\ \sum_{i,j} \pi_i P_{i,j} d_{i,j}^k \leq V^k, \forall k \in K \end{aligned} \quad (16)$$

Examples of ergodic coupling constraints are given in Section 2.3.

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 [31] 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:

$$\liminf_{T \rightarrow +\infty} \frac{1}{T} \mathbb{E} \left(\sum_{t=0}^{T-1} r(X_t, \nu_t) \right) \quad (17)$$

under the $|K|$ *ergodic constraints*

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

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 [31] 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 (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 5), it is easy to see that this linear program is equivalent to:

$$\max_{\rho} \left\{ \sum_{i,j \in [n]} \rho_{i,j} r_{i,j} \text{ st: } \rho \in \mathcal{R} \text{ and } \sum_{i,j \in [n]} \rho_{i,j} d_{i,j}^k \leq V^k, \forall k \in K \right\} \quad (18)$$

where \mathcal{R} is the image of $\prod_{i \in [n]} \mathcal{P}_i$ by the correspondence of Proposition 1. The set \mathcal{R} is a polyhedron, as soon as every \mathcal{P}_i is a polyhedron (Proposition 2).

Following the correspondence discussed in Proposition 1, we can see that the linear Problem (18) is just the reformulation of Problem (16) in terms of occupation measures when we consider total income utility (4).

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

Proposition 11. *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 (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 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 [32] 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 15. Proposition 11 applies in particular if W is a relative entropy utility function, ie $W(\rho) = -\sum_{i,j \in [n]} \rho_{i,j} \log(\rho_{i,j}/\mu_{i,j})$, where parameters $\mu_{i,j} > 0$ (the reference measure) are given.

If we choose to minimize the entropy function on the whole web graph, we recover the TrafficRank algorithm [33]. 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 websurfers on her website, interpreted as a fair distribution of websurfers.

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.

6.2 Optimality condition

The following shows that the mean reward before teleportation (Definition 4) 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 12. *The derivative of total utility function (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 6 (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). 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 introduce the numbers $D_{i,j}^k = \pi_i d_{i,j}^k + \pi d^k (I - \alpha S)^{-1} e_i P_{i,j}$. 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 (2) 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 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 [34]. 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 $T_{\mathcal{P}^d}(P) = \left\{ Q \in \mathbb{R}^{n \times n} \mid \nabla f^{-1} Q \in T_{\mathcal{R}^d}(f^{-1}(P)) \right\}$.

$\nabla f^{-1} Q \in T_{\mathcal{R}^d}(f^{-1}(P))$ first means that $\nabla f^{-1} Q \in T_{\mathcal{R}}(f^{-1}(P))$ which can also be written as $\nabla f \nabla f^{-1} Q = Q \in T_{\mathcal{P}}(P)$. 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 12, we get the expression stated in the theorem.

By Proposition 11, 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

6.3 A Lagrangian relaxation scheme to handle coupling constraints between pages

The PageRank Optimization Problem with "ergodic" coupling constraints (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 13. *The PageRank Optimization problem with K "ergodic" coupling constraints (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 [35] Theorem 21 and Remark 33 for instance. Here we relax the coupling constraints in the problem written with occupation measures (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.

7 Experimental results

7.1 Continuous problem with local constraints only

We have tried our algorithms on a crawl on eight New Zealand Universities available at [36]. 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.

Assume that the webmasters controlling these pages 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 2.3). 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 `maori-oteha.massey.ac.nz/te_waka.htm`. That page should link to `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 of 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.

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 4.3 (60 iterations). We compared our algorithm with an adaptation of the graph augmentation approach of [17] to total utility: 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 962873 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.

7.3 Coupling linear constraints

We would like to solve the discrete optimization problem of the preceding section with two additional coupling 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 2.3). We also require that the sum of PageRank val-

ues of the home pages of the 10 universities considered remains at least equal to their initial value after the optimization (effective frequency constraint).

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 13, 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 0.00769. Computation took 675 s (11 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 PageRank values of 0.00756. Thus we have here a duality gap of at most 1.7%. 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 124,328 internal links but also 11,235 external links. As in Section 7.2, lots of links are added, but here there are also external links.

The bounding technique proposed here can also be adapted to PageRank optimization problem with mutual exclusion constraints. It may also be possible to use it to design a branch and bound algorithm to solve the problem exactly thanks to the bounds found.

Conclusion

We have presented in this paper 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 an other. This means 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 to address more complex problems, for instance with coupling constraints, in order to get less trivial solutions.

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. The latter is the object of further research.

References

- [1] S. Brin and L. Page, “The anatomy of a large-scale hypertextual web search engine,” *Computer Networks and ISDN Systems*, vol. 30, no. 1-7, pp. 107–117, 1998, proc. 17th International World Wide Web Conference.
- [2] D. Chaffey, C. Lake, and A. Friedlein, *Search Engine Optimization - Best Practice Guide*. Econsultancy.com Ltd, 2009.
- [3] Z. Gyöngyi and H. Garcia-Molina, “Link spam alliances,” in *VLDB ’05: Proc. 31st international conference on Very large data bases*. VLDB Endowment, 2005, pp. 517–528.
- [4] R. A. Baeza-Yates, C. Castillo, and V. López, “Pagerank increase under different collusion topologies,” in *Proc. of the First International Workshop on Adversarial Information Retrieval on the Web*, 2005, pp. 17–24.
- [5] P. Berkhin, “A survey on pagerank computing,” *Internet Mathematics*, vol. 2, pp. 73–120, 2005.
- [6] A. N. Langville and C. D. Meyer, *Google’s PageRank and beyond: the science of search engine rankings*. Princeton University Press, 2006.
- [7] A. Arasu, J. Novak, A. Tomkins, and J. 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.

- [8] H. Ishii and R. Tempo, “Distributed randomized algorithms for the pagerank computation,” *Automatic Control, IEEE Transactions on*, vol. 55, no. 9, pp. 1987–2002, 2010.
- [9] A. Nazin and B. 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)*, 2009, pp. 127–132.
- [10] A. Borodin, G. O. Roberts, and J. S. Rosenthal, “Link analysis ranking: algorithms, theory, and experiments,” *ACM Transactions on Internet Technology*, vol. 5, pp. 231–297, 2005.
- [11] A. Bonato and W. Laurier, “A survey of models of the web graph,” in *Proc. of Combinatorial and Algorithmic Aspects of Networking*. Springer, 2004, pp. 159–172.
- [12] K. Avrachenkov and N. Litvak, “The Effect of New Links on Google PageRank,” *Stochastic Models*, vol. 22, no. 2, pp. 319–331, 2006.
- [13] F. Mathieu and L. Viennot, “Local aspects of the global ranking of web pages,” in *6th International Workshop on Innovative Internet Community Systems (I2CS)*, Neuchâtel, 2006, pp. 1–10.
- [14] C. de Kerchove, L. Ninove, and P. van Dooren, “Maximizing pagerank via outlinks,” *Linear Algebra and its Applications*, vol. 429, no. 5-6, pp. 1254–1276, 2008.
- [15] L. Ninove, “Dominant vectors of nonnegative matrices,” PhD Thesis, Université Catholique de Louvain, 2008.
- [16] H. Ishii and R. Tempo, “Computing the pagerank variation for fragile web data,” *SICE J. of Control, Measurement, and System Integration*, vol. 2, no. 1, pp. 1–9, 2009.
- [17] B. C. Csáji, R. M. Jungers, and V. D. Blondel, “Pagerank optimization in polynomial time by stochastic shortest path reformulation,” in *Proc. 21st International Conference on Algorithmic Learning Theory*, ser. Lecture Notes in Computer Science, vol. 6331. Springer, 2010, pp. 89–103.

- [18] —, “Pagerank optimization by edge selection,” 2009, eprint arXiv:0911.2280v1.
- [19] D. Nemirovsky and K. Avrachenkov, “Weighted pagerank: Cluster-related weights,” in *Proc. 17th Text REtrieval Conference, TREC’08*, E. M. Voorhees and L. P. Buckland, Eds. National Institute of Standards and Technology (NIST), 2008.
- [20] C. H. Papadimitriou and J. N. Tsitsiklis, “The complexity of markov decision processes,” *Mathematics of Operations Research*, vol. 12, no. 3, pp. 441–450, 1987.
- [21] M. Groetschel, L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, ser. Algorithms and Combinatorics, R. L. Graham, B. Korte, and L. Lovasz, Eds. Springer-Verlag, 1988, vol. 2.
- [22] R. Givan, S. Leach, and T. Dean, “Bounded-parameter markov decision processes,” *Artificial Intelligence*, vol. 122, no. 1-2, pp. 71–109, 2000.
- [23] H. Xu and S. Mannor, “Distributionally robust markov decision processes,” in *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, Eds., 2010, pp. 2505–2513.
- [24] B. Kveton, M. Hauskrecht, and C. Guestrin, “Solving factored mdps with hybrid state and action variables,” *Journal of Artificial Intelligence Research*, vol. 27, pp. 153–201, 2006.
- [25] H.-F. Winkler and S. Vlasceanu, “Quelle est la valeur d’un hyperlien ?” Rapport d’Enseignement d’Approfondissement en Recherche Opérationnelle, Ecole Polytechnique, 2007.
- [26] W. T. Tutte, *Graph theory*, ser. Enc. of Mathematics and its Applications. Cambridge University Press, 2001, vol. 21.
- [27] G. M. Ziegler, *Lectures on Polytopes (Graduate Texts in Mathematics)*. Springer, 2001.
- [28] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.
- [29] “The stanford webbase project.” [Online]. Available: <http://diglib.stanford.edu:8091/~testbed/doc2/WebBase/>

- [30] F. Mathieu, “Graphes du web, mesures d’importance à la pagerank,” Ph.D. dissertation, Université Montpellier II, 2004.
- [31] E. Altman, *Constrained Markov decision processes*, ser. Stochastic Modeling. Chapman & Hall/CRC, 1999.
- [32] A. Ben-Tal and A. Nemirovski, *Lectures on Modern Convex Optimization, Analysis, Algorithms and Engineering Applications*, ser. MPS/SIAM Series on Optimization. SIAM, 2001.
- [33] J. A. Tomlin, “A new paradigm for ranking pages on the world wide web,” in *Proc. 12th international conference on World Wide Web*, ser. WWW ’03. New York, NY, USA: ACM, 2003, pp. 350–355.
- [34] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*, ser. Grundlehren der math. Wiss. Springer, 1997, vol. 317.
- [35] C. Lemaréchal, “Lagrangian relaxation,” in *Computational Combinatorial Optimization*, M. Jünger and D. Naddef, Eds. Springer Verlag, Heidelberg, 2001, pp. 115–160.
- [36] A. W. L. D. Project, “New zealand university web sites,” January 2006. [Online]. Available: <http://cybermetrics.wlv.ac.uk/database/index.html>