Random Walks with "Back Buttons"

(Extended Abstract)

Ronald Fagin* Anna R. Karlin† Jon Kleinberg‡ Prabhakar Raghavan* Sridhar Rajagopalan*
Ronitt Rubinfeld§ Madhu Sudan¶ Andrew Tomkins*

Abstract

We introduce *backoff processes*, an idealized stochastic model of browsing on the world-wide web, which incorporates both hyperlink traversals and use of the "back button." With some probability the next state is generated by a distribution over out-edges from the current state, as in a traditional Markov chain. With the remaining probability, however, the next state is generated by clicking on the back button, and returning to the state from which the current state was entered by a "forward move". Repeated clicks on the back button require access to increasingly distant history.

We show that this process has fascinating similarities to and differences from Markov chains. In particular, we prove that like Markov chains, backoff processes always have a limit distribution, and we give algorithms to compute this distribution. Unlike Markov chains, the limit distribution may depend on the start state.

1 Introduction

Consider a modification of a Markov chain in which at each step, with some probability, we *undo* the last forward transition of the chain. For intuition, the reader may wish to think of a user using a browser on the world-wide web where he is following a Markov chain on the pages of the web, and occasionally hitting the "back button". We model such phenomena by discrete-time stochastic processes of the following form: we are given a Markov chain M on a set $V = \{1, 2, \ldots, n\}$ of states, together with an n-dimensional vector $\vec{\alpha}$ of *backoff probabilities*. The process evolves as follows: at each time step $t = 0, 1, 2, \ldots$, the process is in a state $X_t \in V$, and in addition has a *history* H_t, which is a stack whose items are states from V. Let $\mathrm{top}(H)$ denote the top of the stack H. At time $t = 0$ the process starts at some state $X_0 \in V$, with the history H_0 containing only the single element X_0. At each subsequent step the process makes either a *forward step* or a *backward step*, by the following rules: (i) if H_t consists of the singleton X_0 it makes a forward step; (ii) otherwise, with probability $\alpha_{\mathrm{top}(H_t)}$ it makes a backward step, and with probability $1 - \alpha_{\mathrm{top}(H_t)}$ it makes a forward step. The forward and backward steps at time t are as follows:

1. In a *forward step*, X_t is distributed according to the successor state of X_{t-1} under M; the state X_t is then pushed onto the history stack H_{t-1} to create H_t.

2. In a *backward step*, the process pops $\mathrm{top}(H_{t-1})$ from H_{t-1} to create H_t; it then moves to $\mathrm{top}(H_t)$ (i.e., the new state X_t equals $\mathrm{top}(H_t)$.)[1]

Under what conditions do such processes have limit distributions, and how such processes differ from traditional Markov chains? We focus in this paper on the time-averaged limit distribution, usually called the "Cesaro limit distribution".[2]

Motivation

Our work is broadly motivated by user modeling for scenarios in which a user with an "undo" capability performs a sequence of actions. A simple concrete setting is that of browsing on the world-wide web. We view the pages of the web as states in a Markov chain, with the transition probabilities denoting the distribution over new pages to which the user can move forward, and the backoff vector denoting for each state the probability that a user enters the state and elects to click the browser's back button rather than continuing to browse forward from that state.

A number of research projects [1, 7, 10] have designed and implemented web intermediaries and learning agents that build simple user models, and used them to personalize the user experience. On the commercial side, user models are exploited to better target advertising on the web based on a user's browsing patterns; see [2] and references therein for theoretical results on these and related

*IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120. {fagin, pragh, sridhar, tomkins}@almaden.ibm.com

†Department of Computer Science, Box 352350, University of Washington, Seattle, WA 98195. karlin@cs.washington.edu

‡Department of Computer Science, Cornell University, Ithaca, NY 14853. kleinber@cs.cornell.edu

§NECI, 4 Independence Way, Princeton, NJ 08540. On leave from Cornell University. ronitt@research.nj.nec.com

¶MIT Laboratory for Computer Science, 545 Technology Square NE43-307, Cambridge, MA 02139, madhu@mit.edu. Supported in part by an MIT-NEC Research Initiation Award, a Sloan Foundation Fellowship and NSF Career Award CCR-9875511.

[1] Note that the condition $X_t = \mathrm{top}(H_t)$ holds for all t, independent of whether the step is a forward step or backward step.

[2] The *Cesaro limit* of a sequence $a_0, a_1, \ldots$ is $\lim_{t\to\infty} \frac{1}{t}\sum_{\tau=0}^{t-1} a_\tau$, if the limit exists. For example, the sequence 0,1,0,1,... has Cesaro limit $1/2$. The *Cesaro limit distribution* at state i is $\lim_{t\to\infty} \frac{1}{t}\sum_{\tau=0}^{t-1} \Pr[X_t = i]$, if the limit exists. By contrast, the *stationary distribution* at state i is $\lim_{t\to\infty} \Pr[X_t = i]$, if the limit exists. Of course, a stationary distribution is always a Cesaro limit distribution. We shall sometimes refer simply to either a stationary distribution or a Cesaro limit distribution as a *limit distribution*.

problems. Understanding more sophisticated models such as ours is interesting in its own right, but could also lead to better user modeling.

Overview of Results

For the remainder of this paper we assume a finite number of states. For simplicity, we assume also that the underlying Markov chain is *irreducible* (i.e., it is possible, with positive probability, to eventually reach each state from each other state) and aperiodic. In particular, M has a stationary distribution, and not just a Cesaro limit distribution. Since some backoff probability α_i may equal 1, these assumptions do not guarantee that the backoff process is irreducible (or aperiodic). We shall focus our attention on the situation where the backoff process is irreducible.[3]

We now give the reader a preview of some interesting and arguably unexpected phenomena that emerge in such "back-button" random walks. Our primary focus is on the Cesaro limit distribution.

Intuitively, if the history stack H_t grows unboundedly with time, then the process "forgets" the start state X_0 (as happens in a traditional Markov process, where $\vec{\alpha}$ is identically zero). On the other hand, if the elements of $\vec{\alpha}$ are all very close to 1, the reader may envision the process repeatedly "falling back" to the start state X_0, so that H_t does not tend to grow unboundedly. What happens between these extremes?

One of our main results is that there is always a Cesaro limit distribution, although there may not be a stationary distribution, even if the backoff process is aperiodic. Consider first the case when all entries of $\vec{\alpha}$ are equal, so that there is a single backoff probability α that is independent of the state. In this case we give a remarkably simple characterization of the limit distribution provided $\alpha < 1/2$: the history grows unboundedly with time, and the limit distribution of the process converges to that of the underlying Markov chain M.

On the other hand, if $\alpha > 1/2$ then the process returns to the start state X_0 infinitely often, the expected history length is finite, and the limit distribution differs in general from that of M, and depends on the start state X_0. Thus, unlike ergodic Markov chains, the limit distribution depends on the starting state.

More generally, consider starting the backoff process in a probability distribution over the states of M; then the limit distribution depends on this initial distribution. As the initial distribution varies over the unit simplex, the set of limit distributions forms a simplex. As α converges to $1/2$ from above, these simplices converge to a single point, which is the limit distribution of the underlying Markov chain.

The transition case $\alpha = 1/2$ is fascinating: the process returns to the start state infinitely often, but the history grows with time and the distribution of the process reaches the stationary distribution of M. These results are described in Section 3.

We have distinguished three cases: $\alpha < 1/2$, $\alpha = 1/2$, and $\alpha > 1/2$. In Section 4, we show that these three cases can be generalized to backoff probabilities that vary from state to state. The generalization depends on whether a certain infinite Markov process (whose states correspond to possible histories) is transient, null, or ergodic respectively (see Section 4 for definitions). It is intuitively clear in the constant α case, for example, that when $\alpha < 1/2$, the history will grow unboundedly. But what happens when some states have backoff probabilities greater than 1/2 and others have backoff probabilities less than 1/2? When does the history grow, and how does the limit distribution depend on M and $\vec{\alpha}$? Even when all the backoff probabilities are less than $1/2$, why should there be a limit distribution?

[3] We would like to make the simplifying assumptions that no α_i equals 1, and that the backoff process is irreducible, but we cannot, since later we are forced to deal with cases where these assumptions do not hold.

We resolve these questions by showing that there exists a potential function of the history that is expected to grow in the transient case (where the history grows unboundedly), is expected to shrink in the ergodic case (where the expected size of the history stack remains bounded), and is expected to remain constant if the process is null. The potential function is a bounded difference martingale, which allows us to use martingale tail inequalities to prove these equivalences. Somewhat surprisingly, we can use this relatively simple characterization of the backoff process to obtain an efficient algorithm to decide, given M and α, whether or not the given process is transient, null or ergodic. We show that in all cases the process attains a Cesaro limit distribution (though the proofs are quite different for the different cases). We also give algorithms to compute the limit probabilities. If the process is either ergodic or null then the limit probabilities are computed exactly by solving certain systems of linear inequalities. However, if the process is transient, then the limit probabilities need not be rational numbers, even if all entries of M and $\vec{\alpha}$ are rational. We show that in this case, the limit probabilities can be obtained by solving a linear system, where the entries of the linear system are themselves the solution to a semidefinite program. This gives us an algorithm to approximate the limit probability vector.

2 Definitions and Notation

We use $(M, \vec{\alpha}, i)$ to denote the backoff process on an underlying Markov chain M, with backoff vector $\vec{\alpha}$, starting from state i. This process is an (infinite) Markov chain on the space of all histories. Formally, a *history stack* (or *history*) $\bar{\sigma}$ is a sequence $\langle \sigma_0, \sigma_1, \ldots, \sigma_l \rangle$ of states of V, for $l \geq 0$. For a history $\bar{\sigma} = \langle \sigma_0, \sigma_1, \ldots, \sigma_l \rangle$, its *length*, denoted $\ell(\bar{\sigma})$, is l (we do not count the start state σ_0 in the length, since it is special). If $\ell(\bar{\sigma}) = 0$, then we say that $\bar{\sigma}$ is an *initial* history. For a history $\bar{\sigma} = \langle \sigma_0, \sigma_1, \ldots, \sigma_l \rangle$, its *top*, denoted $\mathrm{top}(\bar{\sigma})$, is σ_l. We also associate the standard stack operations pop and push with histories. For a history $\bar{\sigma} = \langle \sigma_0, \sigma_1, \ldots, \sigma_l \rangle$, we have $\mathsf{pop}(\bar{\sigma}) = \langle \sigma_0, \sigma_1, \ldots, \sigma_{l-1} \rangle$, and for state $j \in \{1, \ldots, n\}$, we have $\mathsf{push}(\bar{\sigma}, j) = \langle \sigma_0, \sigma_1, \ldots, \sigma_l, j \rangle$. We let $\mathcal{S}$ denote the space of all finite attainable histories.

For a Markov chain M, backoff vector $\vec{\alpha}$, and history $\bar{\sigma}$ with $\mathrm{top}(\bar{\sigma}) = j$, define the successor (or next state) $\mathrm{succ}(\bar{\sigma})$ to take on values from $\mathcal{S}$ with the following distribution:

If $\ell(\bar{\sigma}) = 0$ then with probability M_{jk}, $\mathrm{succ}(\bar{\sigma}) = \mathsf{push}(\bar{\sigma}, k)$. If $\ell(\bar{\sigma}) \geq 1$ then

$$\mathrm{succ}(\bar{\sigma}) = \begin{cases} \mathsf{pop}(\bar{\sigma}) & \text{with prob } \alpha_j \\ \mathsf{push}(\bar{\sigma}, k) & \text{with prob } (1-\alpha_j)M_{jk} \end{cases}$$

For a Markov chain M, backoff vector $\vec{\alpha}$ and state $i \in \{1, \ldots, n\}$, the $(M, \vec{\alpha}, i)$-*Markov chain* is the sequence $\langle H_0, H_1, H_2, \ldots \rangle$ taking values from the set $\mathcal{S}$ of histories, with $H_0 = \langle i \rangle$ and H_{t+1} distributed as $\mathrm{succ}(H_t)$. We refer to the sequence $\langle X_0, X_1, X_2, \ldots \rangle$, with $X_t = \mathrm{top}(H_t)$ as the $(M, \vec{\alpha}, i)$-*backoff process*. Several properties of the $(M, \vec{\alpha}, i)$-backoff process are actually independent of the start state i, and to stress this aspect we will sometimes use simply the term "$(M, \vec{\alpha})$-backoff process".

Note that the $(M, \vec{\alpha}, i)$-backoff process does not completely give the $(M, \vec{\alpha}, i)$-Markov chain, because it does not specify whether each step results from a "forward" or "backward" operation. To complete the correspondence we define an *auxiliary sequence*: Let $S_1, \ldots, S_t, \ldots$ be the sequence with S_t taking on values from the set $\{F, B\}$, with $S_t = F$ if $\ell(H_t) = \ell(H_{t-1}) + 1$ and $S_t = B$ if $\ell(H_t) = \ell(H_{t-1}) - 1$. (Intuitively, F stands for "forward" and B for "backward".) Notice that sequence $X_0, \ldots, X_t, \ldots$ together with the sequence $S_1, \ldots, S_t, \ldots$ does completely specify the sequence $H_0, \ldots, H_t, \ldots$.

3 Constant Backoff Probability

The case in which the backoff probability takes the same value α for every state has a very clean characterization, and it will give us insight into some of the arguments to come. In this case, we refer to the $(M, \vec{\alpha}, i)$-backoff process as the (M, α, i)-backoff process (where we drop the vector sign above α).

We fix a specific (M, α, i)-backoff process throughout this section. Suppose we generate a sequence $X_0, X_1, \ldots, X_t, \ldots$ of steps together with an auxiliary sequence $S_1, \ldots, S_t, \ldots$. To begin with, we wish to view this sequence of steps as being "equivalent" (in a sense) to one in which only forward steps are taken. In this way, we can relate the behavior of the (M, α, i)-backoff process to that of the underlying (finite) Markov process M beginning in state i, which we understand much more accurately. We write $q_t(j)$ to denote the probability that M, starting in state i, is in state j after t steps.

When the backoff probability takes the same value α for every state, we have the following basic relation between these two processes.

Theorem 3.1 *For given natural numbers λ and t, and a state j, we have* $\Pr[X_t = j \mid \ell(H_t) = \lambda] = q_\lambda(j)$.

We omit the proof of this theorem due to space constraints.

In addition to the sequences $\{X_t\}$ and $\{S_t\}$, consider the sequence $\{Y_t : t \geq 0\}$, where Y_t is the history length $\ell(H_t)$. Now Y_t is simply the position after t steps of a random walk on the natural numbers, with a reflecting barrier at 0, in which the probability of moving left is α and the probability of moving right is $1 - \alpha$. This correspondence will be crucial for our analysis.

In terms of these notions, we mention one additional technical lemma. Its proof follows simply by conditioning on the value of Y_t and applying Theorem 3.1.

Lemma 3.2 $\Pr[X_t = j] = \sum_r q_r(j) \cdot \Pr[Y_t = r]$, *for every natural number t and state j.*

We are now ready to consider the two cases where $\alpha \leq \frac{1}{2}$ and where $\alpha > \frac{1}{2}$, and show that in each case there is a Cesaro limit distribution.

The Case of Backoff Probability at Most 1/2:

Let the stationary probability distribution of the underlying Markov chain M be $\langle \psi_1, \ldots, \psi_n \rangle$. By our assumptions about M, this distribution is independent of the start state i. When $\alpha \leq \frac{1}{2}$, we show that the (M, α, i)-backoff process converges to $\langle \psi_1, \ldots, \psi_n \rangle$. That is, there is a stationary probability distribution, which is independent of the start state i, and this stationary probability distribution equals the stationary probability distribution of the underlying Markov chain.

Theorem 3.3 *For all states j of the (M, α, i)-backoff process, we have* $\lim_{t \to \infty} \Pr[X_t = j] = \psi_j$.

Proof: Fix $\epsilon > 0$, and choose t_0 large enough that for all states j of M and all $t \geq t_0$, we have $|q_t(j) - \psi_j| < \epsilon/2$. Since $\alpha \leq 1/2$, we can also choose $t_1 \geq t_0$ large enough that for each $t \geq t_1$, we have $\Pr[Y_t > t_0] > 1 - \epsilon/2$. Then for $t \geq t_1$ we have

$$\begin{aligned}
&|\Pr[X_t = j] - \psi_j| \\
&= \left| \sum_r q_r(j) \cdot \Pr[Y_t = r] - \psi_j \sum_r \Pr[Y_t = r] \right| \\
&= \left| \sum_r (q_r(j) - \psi_j) \cdot \Pr[Y_t = r] \right| \\
&\leq \sum_r |q_r(j) - \psi_j| \cdot \Pr[Y_t = r] \\
&= \sum_{r < t_1} |q_r(j) - \psi_j| \cdot \Pr[Y_t = r] \\
&\quad + \sum_{r \geq t_1} |q_r(j) - \psi_j| \cdot \Pr[Y_t = r] \\
&\leq \sum_{r < t_1} \Pr[Y_t = r] + \sum_{r \geq t_1} \epsilon/2 \cdot \Pr[Y_t = r] \\
&\leq \epsilon/2 + \epsilon/2 = \epsilon.
\end{aligned}$$

■

Although the proof above applies to each $\alpha \leq \frac{1}{2}$, we note a qualitative difference between the case of $\alpha < \frac{1}{2}$ and the "threshold case" $\alpha = \frac{1}{2}$. In the former case, for every r, there are almost surely only finitely many t for which $Y_t \leq r$; the largest such t is a step on which the process pushes a state that is never popped in the future. In the latter case, Y_t almost surely returns to 0 infinitely often, and yet the process still converges to the stationary distribution of M.

The Case of Backoff Probability Greater than 1/2:

When $\alpha > \frac{1}{2}$, the (M, α, i)-backoff process retains positive probability on short histories as t increases, and hence retains memory of its start state i. Nevertheless, the process has a Cesaro limit distribution, but this distribution may be different from the stationary distribution of M.

Theorem 3.4 *When $\alpha > \frac{1}{2}$, the (M, α, i)-backoff process has a Cesaro limit distribution.*

Proof: Let t be a natural number and j a state. Then $\Pr[X_t = j] = \sum_r q_r(j) \cdot \Pr[Y_t = r]$ by Lemma 3.2. Viewing Y_t as a random walk on the natural numbers, one can compute the Cesaro limit of $\Pr[Y_t = r]$ to be $\zeta_r = \beta\alpha$ when $r = 0$, and $\zeta_r = \beta z^{r-1}$ when $r > 0$, where $\beta = (2\alpha - 1)/(2\alpha^2)$ and $z = (1 - \alpha)/\alpha$. (Note that Y_t does not have a stationary distribution, because it is even only on even steps.) A standard argument then shows that $\Pr[X_t = j]$ has the Cesaro limit $\sum_r \zeta_r q_r(j)$. ■

Note that the proof shows only a Cesaro limit distribution, rather than a stationary distribution. In fact, it is not hard to show that if $\alpha > \frac{1}{2}$, then there is not necessarily a stationary distribution, even if the backoff process is aperiodic.

Now, more generally, suppose that the process starts from an initial distribution over states; we are given a probability vector $z = \langle z_1, \ldots, z_n \rangle$, choose a state j with probability z_j, and begin the process from j. As z ranges over all possible probability vectors, what are the possible vectors of limit distributions? Let us again assume a fixed underlying Markov chain M, and denote this set of limit distributions by S_α. We obtain the following theorem whose proof is omitted for lack of space.

Theorem 3.5 *Each S_α is a simplex. As α converges to $\frac{1}{2}$ from above, these simplices converge to the single vector that is the stationary distribution of the underlying Markov chain.*

4 Varying Backoff Probabilities

Recall that the state space $\mathcal{S}$ of the $(M, \vec{\alpha}, i)$-Markov chain contains all finite attainable histories of the backoff process. Let us refer to the transition probability matrix of the $(M, \vec{\alpha}, i)$-Markov chain as the *Polish matrix with starting state i*, or simply the *Polish matrix* if i is implicit or irrelevant. Note that even though the backoff process has only finitely many states, the Polish matrix has a countably infinite number of states.

Our analysis in the rest of the paper will branch, depending on whether the Polish matrix is transient, null, or ergodic. We now define these concepts, which are standard notions in the study of denumerable Markov chains (see e.g. [6]). A Markov chain is called *recurrent* if, started in an arbitrary state i, the probability of eventually returning to state i is 1. Otherwise, it is called *transient*. There are two subcases of the recurrent case. If, started in an arbitrary state i, the expected time to return to i is finite, then the Markov chain is called *ergodic*. If, started in an arbitrary state i, the probability of return to state i is 1, but the expected time to return to i is infinite, then the Markov chain is called *null*. Every irreducible Markov chain is either transient, ergodic, or null, and for irreducible Markov chains, we can replace every occurrence of "an arbitrary state" by "some state" in these definitions above. Every irreducible Markov chain with a finite state space is ergodic.

As examples, consider a random walk on the semi-infinite line, with a reflecting barrier at 0, where the probability of moving left (except at 0) is p, of moving right (except at 0) is $1 - p$, and of moving right at 0 is 1. If $p < 1/2$, then the walk is transient; if $p = 1/2$, then the walk is null; and if $p > 1/2$, then the walk is ergodic.

We say that the backoff process $(M, \vec{\alpha}, i)$ is transient (resp., null, ergodic) if the Polish matrix is transient (resp., null, ergodic). In the constant α case (Section 3), if $\alpha < 1/2$, then the backoff process is transient; if $\alpha = 1/2$, then the backoff process is null; and if $\alpha > 1/2$, then the backoff process is ergodic. The next proposition says that the classification does not depend on the start state and therefore we may refer to the backoff process $(M, \vec{\alpha})$ as being transient, ergodic, or null. Its proof is omitted for lack of space.

Proposition 4.1 *The irreducible backoff process $(M, \vec{\alpha}, i)$ is transient (resp., ergodic, null) precisely if the backoff process $(M, \vec{\alpha}, j)$ is transient (resp., ergodic, null).*

In the rest of this section, we will outline the proofs of the following three theorems.

Theorem 4.2 *If $(M, \vec{\alpha})$ is irreducible, then the task of classifying the $(M, \vec{\alpha})$-backoff process as transient or ergodic or null is solvable in polynomial time.*

Theorem 4.3 *Each $(M, \vec{\alpha}, i)$-backoff process has a Cesaro limit distribution. If the process is irreducible and is either transient or null, then this limit distribution is independent of i. Furthermore, the limit distribution is computable in polynomial time if the process is ergodic or null.*

When the $(M, \vec{\alpha}, i)$-backoff process is transient, the limit probabilities are not necessarily rational in the entries of M and $\vec{\alpha}$ (see example in Section 4.2.3) and therefore we cannot hope to compute them exactly. Instead, we give an algorithm for approximating these limit probabilities. Specifically, we show the following:

Theorem 4.4 *Let $(M, \vec{\alpha}, i)$ be a transient backoff process on n states, and let all entries of M and $\vec{\alpha}$ be rationals expressible as ratios of l-bit integers. Then given any error bound $\epsilon > 0$, a vector π' that ϵ-approximates the limit distribution π (i.e., satisfies $|\pi'_j - \pi_j| \le \epsilon$) can be computed in time polynomial in n, l and $\log \frac{1}{\epsilon}$.*

A convenient tool in this section is the Perron-Frobenius Theorem, which is as follows:

Theorem 4.5 *(Perron-Frobenius Theorem; see e.g. [4, p. 508]) Let A be an irreducible, nonnegative square matrix. Then*

- *there exists $\vec{v}$, with all components positive, and $\lambda_0 > 0$ such that $A\vec{v} = \lambda_0 \vec{v}$;*
- *if $\lambda \ne \lambda_0$ is any other eigenvalue of A, then $|\lambda| < \lambda_0$;*
- *each $\vec{w}$ such that $A\vec{w} = \lambda_0 \vec{w}$ is a constant multiple of $\vec{v}$; and*
- *each nonnegative eigenvector of A is a constant multiple of $\vec{v}$.*

4.1 Classifying the Backoff Process

In this section we show how it is possible to classify, in polynomial time, the behavior of each irreducible $(M, \vec{\alpha})$-backoff process as transient or ergodic or null. In Section 3 (where the backoff probability is independent of the state), except for initial histories the expected length of the history either always grows, always shrinks, or always stays the same, independent of the top state in the history stack. To see that this argument cannot carry over to this section, consider a simple Markov chain M on two states 1 and 2, with $M_{ij} = 1/2$ for every pair i, j, and with $\alpha_1 = .99$ and $\alpha_2 = .01$. It is clear that if the top state is 1, then the history is expected to shrink while if the top state is 2, then the history is expected to grow. To deal with this imbalance between the states, we associate a weight w_i with every state i and consider the weighted sum of states on the stack. Our goal is to find a weight vector with the property that the sum of the weights of the states in the stack is expected to grow (resp., shrink, remain constant) if and only if the length of the history is expected to grow (resp., shrink, remain constant) This hope motivates our next few definitions.

Definition 4.6 *For a nonnegative vector $\vec{w} = \langle w_1, \ldots, w_n \rangle$ and a history $\bar{\sigma} = \langle \sigma_0, \ldots, \sigma_l \rangle$ of a backoff process on n states, define the $\vec{w}$-potential of $\bar{\sigma}$, denoted $\Phi_{\vec{w}}(\bar{\sigma})$, to be $\sum_{i=1}^{l} w_{\sigma_i}$ (i.e., the sum of the weights of the states in the history, except the start state).*

Definition 4.7 *For a nonnegative vector $\vec{w} = \langle w_1, \ldots, w_n \rangle$ and a history $\bar{\sigma} = \langle \sigma_0, \ldots, \sigma_l \rangle$ of a backoff process on n states, define the $\vec{w}$-differential of $\bar{\sigma}$, denoted $\Delta\Phi_{\vec{w}}(\bar{\sigma})$, to be $\mathrm{E}\left[\Phi_{\vec{w}}(\mathrm{succ}(\bar{\sigma}))\right] - \Phi_{\vec{w}}(\bar{\sigma})$. (Here E represents the expected value over the distribution given by $\mathrm{succ}(\bar{\sigma})$.)*

The following proposition is immediate from the definition.

Proposition 4.8 *If $\bar{\sigma}$ and $\bar{\sigma}'$ are non-initial histories with the same top state j, then*

$$\Delta\Phi_{\vec{w}}(\bar{\sigma}) = \Delta\Phi_{\vec{w}}(\bar{\sigma}') = -\alpha_j w_j + (1 - \alpha_j) \sum_{k=1}^{n} M_{jk} w_k.$$

The above proposition motivates the following definition.

Definition 4.9 *For a nonnegative vector $\vec{w} = \langle w_1, \ldots, w_n \rangle$, a history $\bar{\sigma} = \langle \sigma_0, \ldots, \sigma_l \rangle$ of a backoff process on n states, and state $j \in \{1, \ldots, n\}$, let $\Delta\Phi_{\vec{w},j} = \Delta\Phi_{\vec{w}}(\bar{\sigma})$, where $\bar{\sigma}$ is any history with $j = \mathrm{top}(\bar{\sigma})$ and $\ell(\bar{\sigma}) > 0$. Let $\Delta\Phi_{\vec{w}}$ denote the vector $\langle \Delta\Phi_{\vec{w},1}, \ldots, \Delta\Phi_{\vec{w},n} \rangle$.*

For intuition, consider the constant α case with weight $w_i = 1$ for each state i. In this case $\Phi_{\vec{w}}(\bar{\sigma})$, the $\vec{w}$-potential of $\bar{\sigma}$, is precisely $\ell(\bar{\sigma})$, and $\Delta\Phi_{\vec{w}}(\bar{\sigma})$, the $\vec{w}$-differential of $\bar{\sigma}$, is the expected change in the size of the stack, which is $1 - 2\alpha$. When $\alpha < 1/2$ (resp., $\alpha = 1/2$, $\alpha > 1/2$), so that the expected change in the size of the stack is positive (resp., 0, negative), the process is transient (resp., null, ergodic).

Similarly, in the varying α case we would like to associate a positive weight with every state so that (1) the expected change in potential in every step has the same sign independent of the top state (i.e., $\vec{w}$ is positive and $\Delta\Phi_{\vec{w}}$ is either all positive or all zero or all negative), and (2) this sign can be used to categorize the process as either transient, null or ergodic precisely as it did in the constant α case.

In general, this will not be possible, say, if some $\alpha_i = 1$ and some other $\alpha_j = 0$. Therefore, we relax this requirement slightly and define the notion of an "admissible" vector (applicable to both the vector of weights and also the vector of changes in potential).

Definition 4.10 *We say that an n-dimensional vector $\vec{v}$ is admissible for a vector $\vec{\alpha}$ if $\vec{v}$ is nonnegative and $v_i = 0$ only if $\alpha_i = 1$. (We will say simply "admissible" instead of "admissible for $\vec{\alpha}$" if $\vec{\alpha}$ is fixed or understood.)*

With this definition in hand, we can prove the following lemma.

Lemma 4.11 *For an irreducible $(M, \vec{\alpha})$-backoff process:*

1. *If there exists an admissible $\vec{w}$ s.t $\Delta\Phi_{\vec{w}}$ is also admissible, then the $(M, \vec{\alpha})$-backoff process is transient.*

2. *If there exists an admissible $\vec{w}$ s.t $-\Delta\Phi_{\vec{w}}$ is also admissible, then the $(M, \vec{\alpha})$-backoff process is ergodic.*

3. *If there exists an admissible $\vec{w}$ s.t $\Delta\Phi_{\vec{w}} = \vec{0}$ then the $(M, \vec{\alpha})$-backoff process is null.*

The proof of this lemma is deferred to the full paper. Roughly speaking though, the idea is to show that $\Phi_{\vec{w}}(\bar{\sigma})$ is a bounded-difference martingale. This enables us to use martingale tail inequalities (e.g., [9, p. 92]) to analyze the long-term behavior of the process.

This explains what could happen if we are lucky with the choice of $\vec{w}$. It does not explain how to find $\vec{w}$, or even why the three cases above are exhaustive. In the rest of this section, we show that the cases are indeed exhaustive and give a efficient algorithm to compute $\vec{w}$. This part of the argument relies on the surprising properties of an $n \times n$ matrix related to the $(M, \vec{\alpha})$-process. We now define this matrix, that we call the *Hungarian matrix*.

Let A be the $n \times n$ diagonal matrix with the ith diagonal entry being α_i. Let I be the $n \times n$ identity matrix. If $\alpha_i > 0$ for every i, then the Hungarian matrix for the $(M, \vec{\alpha})$-process, denoted $H = H_{(M,\vec{\alpha})}$ is the matrix $(I - A)MA^{-1}$. (Notice that A^{-1} does exist and is the diagonal matrix with ith entry being $1/\alpha_i$.)

The spectral properties of H, and in particular its maximal eigenvalue, denoted $\rho(H)$, play a central role in determining the behavior of the $(M, \vec{\alpha})$-process. In this section we show how it determines whether the process is ergodic, null, or transient. In later sections, we will use it to compute limit probability vectors, for a given $(M, \vec{\alpha})$-process.

The maximal eigenvalue $\rho(H)$ motivates us to define a quantity $\rho(M, \vec{\alpha})$ which is essentially equal to $\rho(H)$, in cases where H is defined.

Definition 4.12 *We define $\rho(M, \vec{\alpha})$ to be the supremum over ρ such that there exists an admissible $\vec{w}$ such that the vector $(I - A)M\vec{w} - \rho A\vec{w}$ is admissible.*

We first dispense with the case where some $\alpha_i = 0$.

Claim 4.13 *If $(M, \vec{\alpha})$ is irreducible and $\alpha_j = 0$ for some j, then $\rho(M, \vec{\alpha}) = \infty$.*

Remark: From the proof it follows that if every entry of M and $\vec{\alpha}$ is an l-bit rational, then for any $\rho \leq 2^l$, there exists a non-negative vector $\vec{w}$ with $\|w\|_\infty \leq 1$ and $w_i \geq 2^{-\mathrm{poly}(n,l)}$ if $w_i \neq 0$ satisfying $(I - A)M\vec{w} \geq \rho\vec{w}$. This fact will be used in Section 4.2.3.

Proof: Let $\rho < \infty$ be any constant. We prove the claim by explicitly constructing an admissible vector $\vec{w}$ such that $(I-A)M\vec{w} - \rho A\vec{w}$ is admissible.

Let $M_{\min}$ be the smallest non-zero entry of M, and let $\alpha_{\max}$ be the largest entry of $\vec{\alpha}$ that is strictly smaller than 1. Let γ be any positive number less than $(1 - \alpha_{\max})M_{\min}/\rho$. Let j be any index s.t. $\alpha_j = 0$. Let $G_{M,\vec{\alpha}}$ be the graph on vertex set $\{1, \ldots, n\}$ that has an edge from i to k, if $\alpha_i \neq 1$ and $M_{ik} \neq 0$. (This is the graph with edges giving forward steps of positive probability of the $(M, \vec{\alpha})$-process.) Let $d(i, k)$ denote the length of the shortest path from i to k in the graph $G_{M,\vec{\alpha}}$. By the irreducibility of the $(M, \vec{\alpha})$-process we have that $d(i, j) < n$ for every state i. We now define $\vec{w}$ as follows.

$$w_i = \begin{cases} 0 & \text{if } \alpha_i = 1 \\ \gamma^{d(i,j)} & \text{otherwise.} \end{cases}$$

It is clear by construction that $\gamma > 0$ and thus $\vec{w}$ is admissible. Let $\vec{v} = (I - A)M\vec{w} - \rho A\vec{w}$. We argue that $\vec{v}$ is admissible componentwise, showing that v_i satisfies the condition of admissibility for every i.

Case 1: $\alpha_i = 1$. In this case it suffices to show $v_i \geq 0$. This follows from the facts that $\sum_k (1 - \alpha_i)M_{ik}w_k \geq 0$, and $-\rho\alpha_i w_i = 0$ since $w_i = 0$.

Case 2: $\alpha_i = 0$. (This includes the case $k = j$.) In this case, again we have $-\rho\alpha_i w_i = 0$. Further we have $\sum_k (1 - \alpha_i)M_{ik} = \sum_k M_{ik} = 1$ and thus $v_i = 1$, which also satisfies the condition for admissibility.

Case 3: $0 < \alpha_i < 1$. In particular, $i \neq j$ and $d(i, j) > 0$. Let k be such that $d(k, j) = d(i, j) - 1$ and there is an edge from i to k in $G_{M,\vec{\alpha}}$. We know such a state k exists (by definition of shortest paths). We have:

$$\begin{aligned} v_i &= \sum_{k'} (1 - \alpha_i)M_{ik'}w_{k'} - \rho\alpha_i w_i \\ &\geq (1 - \alpha_i)M_{ik}w_k - \rho\alpha_i w_i \\ &\geq (1 - \alpha_{\max})M_{\min}w_k - \rho w_i \\ &= (1 - \alpha_{\max})M_{\min}\gamma^{d(k,j)} - \rho\gamma^{d(i,j)} \\ &= ((1 - \alpha_{\max})M_{\min} - \rho\gamma)\,\gamma^{d(k,j)} \\ &> 0 \quad (\text{since } \gamma < \tfrac{(1-\alpha_{\max})M_{\min}}{\rho}) \end{aligned}$$

Again the condition for admissibility is satisfied. ■

The next claim shows that in the remaining cases $\rho(M, \vec{\alpha}) = \rho(H)$.

Claim 4.14 *Let $(M, \vec{\alpha})$ be irreducible. If $\alpha_i > 0$ for every i, then $\rho(M, \vec{\alpha}) = \rho(H)$. Further, there exists an admissible vector $\vec{w}$ such that $(I - A)M\vec{w} = \rho(M, \vec{\alpha})A\vec{w}$.*

Proof: Note first that the Hungarian matrix H is nonnegative. Our hope is to apply the Perron-Frobenius theorem (Theorem 4.5) to this non-negative matrix and derive some benefits from this. However, H is not necessarily irreducible, so we can not do this yet. So we consider a smaller matrix, $H|_{\vec{\alpha}}$, which is the restriction of H to rows and columns corresponding to j such that $\alpha_j < 1$. Notice that $H|_{\vec{\alpha}}$ is irreducible. (This is equivalent to $M|_{\vec{\alpha}}$ being irreducible, which is implied by the irreducibility of the backoff process.) By Theorem 4.5, there exists a (unique) positive vector $\vec{v}'$ and a (unique) positive real $\rho = \rho(H|_{\vec{\alpha}})$ such that $H|_{\vec{\alpha}}\vec{v}' = \rho\vec{v}'$. In what follows we see that $\rho(M, \vec{\alpha}) = \rho(H|_{\vec{\alpha}}) = \rho(H)$.

First we verify that $\rho(H|_{\vec{\alpha}}) = \rho(H)$. This is easily seen to be true. Note that the rows of H that are omitted from $H|_{\vec{\alpha}}$ are all 0. Thus a vector $\vec{x}$ is a right eigenvector of H if and only if it is obtained from a right eigenvector $\vec{x}'$ of $H|_{\vec{\alpha}}$ by padding with zeroes (in indices j where $\alpha_j = 1$), and this padding preserves eigenvalues. In particular, we get that $\rho(H) = \rho(H|_{\vec{\alpha}})$ and there is an admissible vector $\vec{v}$ (obtained by padding $\vec{v}'$) such that $H\vec{v} = \rho(H)\vec{v}$.

Next we show that $\rho(M, \vec{\alpha}) \geq \rho(H)$. Consider any $\rho' < \rho(H)$ and let $\vec{w} = A^{-1}\vec{v}$. Then note that $(I - A)M\vec{w} - \rho' A\vec{w} = H\vec{v} - \rho'\vec{v} = (\rho(H) - \rho')\vec{v}$ which is admissible. Thus $\rho(M, \vec{\alpha}) \geq \rho'$ for every $\rho' < \rho(H)$ and thus $\rho(M, \vec{\alpha}) \geq \rho(H)$.

Finally we show that $\rho(M, \vec{\alpha}) \leq \rho(H)$. Let $\vec{w}$ be an admissible vector and let $\rho > 0$ be such that $(I - A)M\vec{w} - \rho A\vec{w}$ is admissible. Let $\vec{v} = A^{-1}\vec{w}$. First note that v_j must be 0 if $\alpha_j = 1$, or else the jth component of the vector $(I - A)M\vec{w} - \rho A\vec{w}$ is negative. Now let $\vec{v}'$ be obtained by restricting $\vec{v}$ to coordinates such that $\alpha_j < 1$. Notice now that we have $H|_{\vec{\alpha}}\vec{v}' - \rho\vec{v}'$ is a non-negative vector. From the fact [8, p. 17] that

$$\rho(A) = \max_{\vec{x}} \left\{ \min_{i|x_i \neq 0} \left\{ \frac{(Ax)_i}{x_i} \right\} \right\}$$

for any irreducible non-negative matrix A, we conclude that $\rho(H|_{\vec{\alpha}}) \geq \rho$.

This concludes the proof that $\rho(M, \vec{\alpha}) = \rho(H)$. The existence of a vector $\vec{w}$ satisfying $(I - A)M\vec{w} = \rho(H)A\vec{w}$ also follows from the argument above. ■

Lemma 4.15 *For every irreducible $(M, \vec{\alpha})$-backoff process, $\rho(M, \vec{\alpha})$ is computable in polynomial time. Furthermore,*

- *$(M, \vec{\alpha})$ is ergodic if and only iff $\rho(M, \vec{\alpha}) < 1$.*
- *$(M, \vec{\alpha})$ is null if and only iff $\rho(M, \vec{\alpha}) = 1$.*
- *$(M, \vec{\alpha})$ is transient if and only iff $\rho(M, \vec{\alpha}) > 1$.*

Proof: The fact that $\rho(M, \vec{\alpha})$ is computable efficiently follows from Claims 4.13 and 4.14.

Notice now that $\Delta\Phi_{\vec{w}} = (I - A)M\vec{w} - A\vec{w}$. We start with the case $\rho(M, \vec{\alpha}) < 1$. Notice that in this case, no $\alpha_i = 0$ (by Claim 4.13) and hence we can apply Claim 4.14 to see that there exists a vector $\vec{w}$ such that $(I - A)M\vec{w} = \rho A\vec{w}$. For this vector $\vec{w}$, we have $\Delta\Phi_{\vec{w}} = (\rho - 1)A\vec{w}$. Thus, the vector $-\Delta\Phi_{\vec{w}} = (1 - \rho)\vec{w}$ is admissible. Applying Lemma 4.11 (part 2), we conclude that the $(M, \vec{\alpha})$-process is ergodic.

Similarly, if $\rho(M, \vec{\alpha}) = 1$, we have that for the vector $\vec{w}$ from Claim 4.14, $\Delta\Phi_{\vec{w}} = \vec{0}$. Thus, by Lemma 4.11 (part 3), we find that the $(M, \vec{\alpha})$-process is null. Finally, if $\rho(M, \vec{\alpha}) > 1$, then (by the definition of $\rho(M, \vec{\alpha})$) there exists a vector $\vec{w}$ and $\rho' > 1$ such that $(I - A)M\vec{w} - \rho' A\vec{w}$ is admissible. In particular, this implies that the vector $\Delta\Phi_{\vec{w}} = (I - A)M\vec{w} - A\vec{w}$ is also admissible. Applying Lemma 4.11 (part 1), we conclude that the $(M, \vec{\alpha})$-process is transient. ■

Note from our proofs that we get the result claimed earlier, that there is a polynomial-time algorithm for computing an admissible $\vec{w}$ such that if the backoff process is transient (resp., ergodic, null), then $\Delta\Phi_{\vec{w}}$ is admissible (resp., $-\Delta\Phi_{\vec{w}}$ is admissible, $\Delta\Phi_{\vec{w}} = \vec{0}$).

Theorem 4.2 follows immediately from Lemma 4.15.

4.2 Cesaro Limit Distributions

We begin this section by sketching the proof that the $(M, \vec{\alpha}, i)$-backoff process always has a Cesaro limit distribution. The proof is quite different in each of the cases (transient, ergodic and null). We conclude the section by showing how the limit distribution may be computed.

The easiest case is the ergodic case. Since the Polish matrix is ergodic, the corresponding Markov process has a Cesaro limit. This gives us a Cesaro limit in the backoff process, where the probability of state i is the sum of the probabilities of the states (stacks) in the Polish matrix with top state i.

We next consider the transient case. When the backoff process is in a state (with a given stack), and that state is never popped off of the stack (by taking a backedge), then we refer to this (occurrence of the) state as *irrevocable*. Let us fix a state i, and consider a renewal process, where each new epoch begins when the process has an irrevocable occurrence of state i. Since the Polish matrix is transient, the expected length of an epoch is finite. The limit probability distribution of state j is the expected number of times that the process is in state j in an epoch, divided by the expected length of an epoch. This is a sketch of a proof of the existence of a Cesaro limit distribution. A more careful argument (omitted here for lack of space) shows the existence of a stationary distribution.

Finally, we consider the null case. We select a state j where $\alpha_j \neq 1$. Let us consider a new backoff process, where the underlying Markov matrix M is the same; where all of the backoff probabilities α_k are the same, except that we change α_j to 1; and where we change the start state to j. This new backoff process can be shown to be ergodic. We are able to show a way of "pasting together" runs of the new ergodic backoff process to simulate runs of the old null ergodic process. Thereby, we show the remarkable fact that the old null process has a Cesaro limit distribution which is the same as the Cesaro limit distribution of the new ergodic process. The details are omitted in this extended abstract.

We now show how the limit distribution may be computed. We can assume without loss of generality that the backoff process is irreducible, since we can easily restrict our attention to an irreducible "component". Again, we branch into three cases. The fact that the limit distribution does not depend on the start state in the null and ergodic cases follows immediately from the fact that our computation in these cases does not depend on the start state.

4.2.1 Null Case

The matrix $H = (I - A)MA^{-1}$, which we saw in Section 4.1), plays an important role in this section. We refer to this matrix as the *Hungarian matrix* of the $(M, \vec{\alpha})$-backoff process. The next theorem gives an important application of the Hungarian matrix.

Theorem 4.16 *The limit probability distribution π satisfies $\pi = \pi H$. This linear system has a unique solution subject to the restriction $\sum_i \pi_i = 1$. Thus, the limit probability distribution can found by solving a linear system.*

Proof: The key ingredient in the proof is the observation that in the null case, the limit probability of a transition from a state i to a state j by a forward step is the same as the limit probability of a transition from state j to a state i by a backward step (since each forward move is eventually revoked, with probability 1). Thus if we let $\pi_{i\to j}$ denote the limit probability of a forward step from i to j and $\pi_{i\leftarrow j}$ denote the limit probability of a backward step from j to i (and π_i denotes the limit probability of being in state i), then the following conditions hold:

$$\pi_i = \sum_j \pi_{i\to j} + \sum_j \pi_{j\leftarrow i}$$

$$\pi_{i\to j} = (1 - \alpha_i)M_{ij}\pi_i; \qquad \pi_{i\to j} = \pi_{i\leftarrow j}.$$

The only controversial condition is the third one, that $\pi_{i\to j} = \pi_{i\leftarrow j}$. The fact that $\pi_{i\leftarrow j}$ exists and equals $\pi_{i\to j}$ appears in the full paper. Manipulating the above conditions shows that π satisfies $\pi = \pi H$.

We now consider uniqueness. Assume first that $\alpha_i < 1$ for every i. Then H is irreducible and nonnegative and thus by the Perron-Frobenius Theorem (Theorem 4.5), it follows easily that π

is the unique solution to the linear system. If some $\alpha_i = 1$, we argue by focusing on the matrix $H|_\alpha$, which is irreducible (as in Section 4.1, $H|_\alpha$ is the principal submatrix of H containing only rows and columns corresponding to i s.t. $\alpha_i < 1$). Renumber the states of M so that the α_i's are non-decreasing. Then the Hungarian matrix looks as follows:

$$H = \begin{bmatrix} H|_\alpha & X \\ 0 & 0 \end{bmatrix},$$

where $H|_\alpha$ is nonnegative and irreducible and X is arbitrary. Write $\pi = [\pi_A \pi_B]$, where π_B has the same number of elements as the number of α_i's that are 1. Then the linear system we have to solve is

$$[\pi_A \pi_B] = [\pi_A \pi_B] \begin{bmatrix} H|_\alpha & X \\ 0 & 0 \end{bmatrix}.$$

This system can be solved by finding $\pi_A = \pi_A H|_\alpha$ and then setting $\pi_B = \pi_A X$. Now π_B is uniquely determined by π_A. Furthermore, π_A is uniquely determined, by the Perron-Frobenius Theorem (Theorem 4.5). This concludes the proof of the theorem. ■

4.2.2 Ergodic Case

In this case also the limit probabilities are obtained by solving linear systems, obtained from a renewal argument. We define "epochs" starting at i by simulating the backoff process as follows. The epoch starts at an initial history with $X_0 = \langle i \rangle$. At the first step the process makes a forward step. At every subsequent unit of time, if the process is back at the initial history, it first flips a coin that comes up B with probability α_i and F otherwise. If the coin comes up B, the end of an epoch is declared.

Notice that the distribution of the length of an epoch starting at i is precisely the same as the distribution of time, starting at an arbitrary non-initial history with i on top of the stack, until this occurrence of i is popped from the stack, conditioned on the fact that the first step taken from i is a forward step.

Let T_i denote the expected length of (or more precisely, number of transitions in) an epoch, when starting at state i. Let N_{ij} denote the expected number of transitions out of state j in an epoch when starting at state i. Standard renewal arguments show that the Cesaro limit probability distribution vector $\pi^{(i)}$, for an $(M, \vec{\alpha}, i)$-backoff process, is given by $\pi_j^{(i)} = N_{ij}/T_i$, provided T_i is finite. This gives us a way to compute the Cesaro limit distribution. The key equations that allow us to compute the N_{ij} and T_i are:

$$T_i = 1 + \sum_k M_{ik}[\alpha_k \cdot 1 + (1-\alpha_k)(T_k+1)] + (1-\alpha_i)T_i, \quad (1)$$

$$\begin{aligned} N_{ij} &= \delta_{ij} + \sum_k M_{ik}[\alpha_k \cdot \delta_{jk} + (1-\alpha_k)(N_{kj} + \delta_{jk})] \\ &\quad +(1-\alpha_i)N_{ij}, \end{aligned} \quad (2)$$

where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise. (The above equations are derived by straightforward conditioning. For example, if the first step in the epoch takes the process to state k, then it takes T_k units of time to return to $\langle i \rangle$ and then with probability $(1 - \alpha_i)$ it takes T_i more steps to end the epoch.)

We claim that the first set (1) of linear equations completely specify T. We argue this as follows. First we may rearrange terms in the equation, and use the fact that $\sum_k M_{ik} = 1$, to simplify (1) to:

$$\alpha_i T_i = 2 + \sum_k (1-\alpha_k) M_{ik} T_k.$$

Dividing both sides by α_i (we know that no $\alpha_i = 0$ in the ergodic case), moving all terms involving T_k to the left, and using the fact that the Hungarian matrix H is given by $H_{ik} = \frac{1-\alpha_k}{\alpha_i} M_{ik}$, we get:

$$T_i - \sum_k H_{ik} T_k = \frac{2}{\alpha_i}.$$

Letting $\vec{T} = \langle T_1, \ldots, T_n \rangle$ and $\vec{b} = \langle 2/\alpha_1, \ldots, 2/\alpha_n \rangle$, we get $(I - H)\vec{T} = \vec{b}$. Since the maximal eigenvalue of H is less than 1, we know that $I - H$ has an inverse (and is given by $I + H + H^2 + \cdots$) and thus $\vec{T}$ is given by $(I-H)^{-1}\vec{b}$.

Similarly, if we let $\vec{N}_j = \langle N_{1j}, \ldots, N_{nj} \rangle$ and $\vec{b}_j = \langle \frac{\delta_{1j} + M_{1j}}{\alpha_1}, \ldots, \frac{\delta_{nj}+M_{nj}}{\alpha_n} \rangle$, then (2) simplifies to yield $\vec{N}_j = (I-H)^{-1}\vec{b}_j$.

Thus $\vec{T}$ and the $\vec{N}_j$'s can be computed using the above linear equations. Using now the formula $\pi_j^{(i)} = N_{ij}/T_i$, we can also compute the stationary probability vectors.

4.2.3 Transient Case

We now prove Theorem 4.4.

Definition 4.17 *For a state j, define the revocation probability as follows: Pick any non-initial history $\bar{\sigma} = \langle \sigma_0, \ldots, \sigma_l \rangle$ with* $\mathrm{top}(\bar{\sigma}) = j$. *The* revocation probability r_j *is the probability that the $(M, \vec{\alpha}, i)$-Markov chain starting at state $\bar{\sigma}$ reaches the state $\bar{\sigma}' = \langle \sigma_0, \ldots, \sigma_{l-1} \rangle$. (Notice that this probability is independent of i, l, and $\sigma_0, \ldots, \sigma_{l-1}$; thus, the quantity is well-defined.)*

Note that r_i is the probability that an epoch starting at i, as in Section 4.2.2, ends in finite time. Let $\vec{r}$ denote the vector of revocation probabilities. The following lemma shows how to compute the limit probabilities π given $\vec{r}$. Further it shows how to compute a close approximation to π, given a sufficiently close approximation to $\vec{r}$.

Lemma 4.18 *The limit probabilities satisfy $\pi = \pi(I - A)MR$, where R is a diagonal matrix with*

$$R_{ii} = \frac{1}{1 - (1-\alpha_i)\sum_k r_k M_{ik}}.$$

Further, there exists a unique solution to the this system subject to the condition $\sum_i \pi_i = 1$.

Remarks: If $\alpha_i = 0$ for every i, then $r_i = 0$ for every i, and so we recover the familiar condition for Markov chains that $\pi = \pi M$. Although we are considering the transient case here, note that if we formally take $r_i = 1$, which occurs in the null case, then we in fact recover the equation we found in the null case, namely $\pi = \pi(I - A)MA^{-1}$.

Proof: The first part of the lemma is obtained as in Theorem 4.16. Let $\pi_{i\to j}$ denote the limit probability of a forward step from i to j, and let $\pi_{i\leftarrow j}$ denote the limit probability of a backward step from j to i. Then the following conditions hold.

$$\pi_{i\leftarrow j} = r_j \pi_{i\to j} \quad (3)$$

$$\pi_{i\to j} = \pi_i (1-\alpha_i) M_{ij} \quad (4)$$

$$\pi_i = \sum_j \pi_{j\to i} + \sum_j \pi_{i\leftarrow j} \quad (5)$$

Using equation (3) to eliminate all occurrences of variables of the form $\pi_{i\leftarrow j}$, and then using equation (4) to eliminate all occurrences of $\pi_{i\to j}$, equation (5) becomes:

$$\pi_i = \sum_j \pi_j (1-\alpha_j) M_{ji} + \pi_i (1-\alpha_i) \sum_j r_j M_{ij} \quad (6)$$

Thus if we let D be the matrix with

$$D_{ij} = \frac{(1-\alpha_i)M_{ij}}{1-(1-\alpha_j)\sum_k M_{jk}r_k},$$

then π satisfies $\pi = \pi D$. As in the proof of Theorem 4.16 if we permute the rows and columns of D so that all states i with $\alpha_i = 1$ appear at the end, then the matrix D looks as follows:

$$D = \left[\begin{array}{cc} D_\alpha & X \\ 0 & 0 \end{array} \right]$$

where D_α is nonnegative and irreducible. Thus $\pi = [\pi_A \pi_B]$ must satisfy $\pi_A = \pi_A D_\alpha$ and $\pi_B = \pi_A X$. Now π_A is seen to be unique (up to scaling) by the Perron-Frobenius Theorem (Theorem 4.5), while π_B is unique given π_A. The lemma follows by noticing that D can be expressed as $(I-A)MR$. ■

Lemma 4.19 *Let the entries of M and $\vec{\alpha}$ be l-bit rationals describing a transient $(M, \vec{\alpha}, i)$-backoff process and let π be its limit probability vector. For every $\epsilon > 0$, there exists $\beta > 0$, with $\log \frac{1}{\beta} = \text{poly}(n, l, \log \frac{1}{\epsilon})$, such that given any vector $\vec{r}'$ of l'-bit rationals satisfying $\|\vec{r}' - \vec{r}\|_\infty \leq \beta$, a vector π' satisfying $\|\pi' - \pi\|_\infty \leq \epsilon$ can be found in time $\text{poly}(n, l, l', \log \frac{1}{\epsilon})$.*

Remark: By truncating $\vec{r}'$ to $\log \frac{1}{\beta}$ bits, we can ensure that l' also grows polynomially in the input size, and thus get a fully polynomial time algorithm to approximate π.

We omit the proof of Lemma 4.19 in this extended abstract.

In the next lemma, we address the issue of how the revocation probabilities may be determined. We show that they form a solution to a quadratic program; in fact a semi-definite program. (Recall that a real symmetric matrix A is positive semidefinite if all of its eigenvalues are non-negative. A semi-definite program is an optimization problem with a linear objective function, whose constraints are of the form "$A[\vec{x}]$ is positive semidefinite", where $A[\vec{x}]$ denotes a symmetric matrix whose entries are themselves linear forms in the variables $x_1, \ldots, x_n$. Semidefinite programs are a special case of convex programs, but more general than linear programs. They can be approximately solved efficiently using the famed ellipsoid algorithm (see [3] for more details).)

Lemma 4.20 *The revocation probabilities r_i are the optimum solution to the following system:*

$$\left. \begin{array}{lrcl} & \min \sum_i x_i & & \\ \textit{s.t.} & x_i & \geq & \alpha_i + (1-\alpha_i)x_i \sum_j M_{ij}x_j \\ & x_i & \leq & 1 \\ & x_i & \geq & 0 \end{array} \right\} \tag{7}$$

Further, the system of inequalities above is equivalent to the following semidefinite program:

$$\left. \begin{array}{lrcl} & \min \sum_i x_i & & \\ \textit{s.t.} & q_i & = & 1-(1-\alpha_i)\sum_j M_{ij}x_j \\ & x_i & \leq & 1 \\ & x_i & \geq & 0 \\ & q_i & \geq & 0 \\ & D_i & & \textit{positive semidefinite,} \end{array} \right\} \tag{8}$$

where

$$D_i = \left[\begin{array}{cc} x_i & \sqrt{\alpha_i} \\ \sqrt{\alpha_i} & q_i \end{array} \right]$$

Proof: We start by considering the following iterative system and proving that it converges to the optimum of (7).

For $t = 0, 1, 2, \ldots$, define $x_i^{(t)}$ as follows:

$$x_i^{(0)} = 0, \quad x_i^{(t+1)} = \alpha_i + (1-\alpha_i)x_i^{(t)} \sum_j M_{ij}x_j^{(t)}.$$

By induction, we note that $x_i^{(t)} \leq x_i^{(t+1)} \leq 1$. The first inequality holds, since

$$\begin{array}{rcl} x_i^{(t+1)} & = & \alpha_i + (1-\alpha_i)x_i^{(t)} \sum_j M_{ij}x_j^{(t)} \\ & \geq & \alpha_i + (1-\alpha_i)x_i^{(t-1)} \sum_j M_{ij}x_j^{(t-1)} \\ & = & x_i^{(t)} \end{array}$$

The second inequality follows similarly. Hence, since $\langle x_i^{(t)} \rangle_t$ is a non-decreasing sequence in the interval $[0, 1]$, it must have a limit. Let x_i^* denote this limit.

We claim that the x_i^* give the (unique) optimum to (7). By construction, it is clear that $0 \leq x_i^* \leq 1$ and $x_i^* = \alpha_i + (1-\alpha_i)x_i^* \sum_j M_{ij}x_j^*$; and hence x_i^*'s form a feasible solution to (7). To prove that it is the optimum, we claim that if $a_1, \ldots, a_n$ are a feasible solution to (7), then we have $a_i \geq x_i^{(t)}$ and thus $a_i \geq x_i^*$. We prove this claim by induction. Assume $a_i \geq x_i^{(t)}$, for every i. Then

$$\begin{array}{rcl} a_i & \geq & \alpha_i + (1-\alpha_i)a_i \sum_j M_{ij}a_j \\ & \geq & \alpha_i + (1-\alpha_i)x_i^{(t)} \sum_j M_{ij}x_j^{(t)} \\ & = & x_i^{(t+1)}. \end{array}$$

This concludes the proof that the x_i^* give the unique optimum to (7).

Next we show that the revocation probability r_i equals x_i^*. To do so, note first that r_i satisfies the condition

$$r_i = \alpha_i + (1-\alpha_i)\sum_j M_{ij}r_j r_i.$$

(Either the move to i is revoked at the first step with probability α_i, or there is a move to j with probability $(1-\alpha_i)M_{ij}$ and then the move to j is eventually revoked with probability r_j, and this places i again at the top of the stack, and with probability r_i this move is revoked eventually.) Thus the r_i's form a feasible solution, and so $r_i \geq x_i^*$. To prove that $r_i \leq x_i^*$, let us define $r_i^{(t)}$ to be the probability that a forward move onto vertex i is revoked in at most t steps. Note that $r_i = \lim_{t \to \infty} r_i^{(t)}$. We will show by induction that $r_i^{(t)} \leq x_i^{(t)}$ and this implies $r_i \leq x_i^*$. Notice first that

$$r_i^{(t+1)} \leq \alpha_i + (1-\alpha_i)\sum_j M_{ij}r_j^{(t)}r_i^{(t)}.$$

(This follows from a conditioning argument similar to the above and then noticing that in order to revoke the move within $t+1$ steps, both the revocation of the move to j and then the eventual revocation of the move to i must occur within t time steps.) Now an inductive argument as earlier shows $r_i^{(t+1)} \leq x_i^{(t+1)}$, as desired. Thus we conclude that $x_i^* = r_i$. This finishes the proof of the first part of the lemma.

For the second part, note that the condition that D_i be semidefinite is equivalent to the condition that $x_i q_i \geq \alpha_i$. Substituting $q_i = 1 - (1-\alpha_i)\sum_j M_{ij} x_j$ turns this into the constraint $x_i - (1-\alpha_i)x_i \sum_j M_{ij}x_j \geq \alpha_i$, and thus establishing the (syntactic) equivalence of (7) and (8). ■

Using Lemmas 4.18 and 4.20 above, we can derive exact expressions for the revocation probabilities and limit probabilities of any given backoff process. The following example illustrates this. It also shows that the limit probabilities are not necessarily rational, even when the entries of M and $\vec{\alpha}$ are rational.

Example: The following example shows that the limit probabilities may be irrational even when all the entries of M and $\vec{\alpha}$ are rational. Let M and $\vec{\alpha}$ be as follows:

$$M = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix} \qquad \vec{\alpha} = \langle \frac{1}{2}, \frac{1}{3} \rangle.$$

Using Lemma 4.20, we can now show that the revocation probabilities are roots of cubic equations. Specifically, r_1 is the unique real root of the equation $-16 + 30x - 13x^2 + 2x^3 = 0$ and r_2 is the unique real root of the equation $-9 + 21x - 14x^2 + 8x^3 = 0$. Both quantities are irrational and given approximately by $r_1 \approx 0.7477$ and $r_2 \approx 0.5775$. Applying Lemma 4.18 to this, we find that the limit probabilities of the $(M, \vec{\alpha})$-process are π_1 and π_2, where π_1 is the unique real root of the equation

$$-1024 + 3936x - 3180x^2 + 997x^3 = 0,$$

and π_2 is the unique real root of the equation

$$-729 + 567x + 189x^2 + 997x^3 = 0.$$

It may be verified that the cubic equations above are irreducible over the rationals, and thus π_1 and π_2 are irrational and given approximately by $\pi_1 \approx 0.3467$ and $\pi_2 \approx 0.6532$.

In the next lemma we show how to efficiently approximate the vector of revocation probabilities. The proof assumes the reader is familiar with standard terminology used in semidefinite programming, and in particular the notion of a separation oracle and its use in the ellipsoid algorithm (see [3] for more details).

Lemma 4.21 *If the entries of M and $\vec{\alpha}$ are given by l-bit rationals, then an ϵ-approximation to the vector of revocation probabilities can be found in time* $\mathrm{poly}(n, l, \log \frac{1}{\epsilon})$.

Proof: We solve the convex program given by (7) approximately using the ellipsoid algorithm [3]. Recall that the ellipsoid algorithm can solve a convex programming problem given (1) a separation oracle describing the convex space, (2) a point $\vec{x}$ inside the convex space, (3) radii ϵ and R such that the ball of radius ϵ around $\vec{x}$ is contained in the convex body and the ball of radius R contains the convex body. The running time is polynomial in the dimension of the space and in $\log \frac{R}{\epsilon}$.

The fact that (7) describes a convex program follows from the fact that it is equivalent to the semidefinite program (8). Further, a separation oracle can also be obtained due to this equivalence. In what follows we will describe a vector $\vec{x}$ that is feasible, and an $\epsilon \geq 2^{-\mathrm{poly}(n,l)}$ such that every point y satisfying $\|x - y\|_\infty \leq \epsilon$ is feasible. Further it is trivial to see that every feasible point satisfies the condition that the ball of radius $\sqrt{n}$ around it contains the unit cube and hence all feasible solutions. This will thus suffice to prove the lemma.

Recall, from Lemma 4.15 of Section 4.1, that since $(M, \vec{\alpha})$ is transient, there exists $\rho > 1$ and a vector $\vec{w}$ such that $(I - A)M\vec{w} \geq \rho A \vec{w}$. Let $w_{\max} = \max_i\{w_i\}$ and $w_{\min} = \min_{i|w_i \neq 0}\{w_i\}$. Notice further that we can choose ρ and $\vec{w}$ such that $\rho \geq 1 + 2^{-\mathrm{poly}(n,l)}$ and $w_{\max} = 1$ and $w_{\min} \geq 2^{-\mathrm{poly}(n,l)}$. (In case $\rho(M, \vec{\alpha}) = \infty$, this follows by picking say $\rho = 2$ and using the remark after Claim 4.13. In case $\rho(M, \vec{\alpha}) < \infty$ we use Claim 4.14 and set $\rho = \rho(H)$ and $\vec{w} = A^{-1}\vec{v}$, where $\vec{v}$ is a right eigenvector of H. Since $\rho > 1$ is an eigenvalue of a matrix whose entries are l-bit rationals and since $\vec{w}$ is a multiple of the eigenvector, the claims about the magnitude of ρ and $w_{\min}$ follow.)

Before describing the vector $\vec{x}$ and ϵ, we make one simplification. Notice that if $\alpha_i = 1$ then $r_i = 1$, and if $\alpha_i = 0$ then $r_i = 0$. We fix this setting and then solve (7) for only the remaining choices of indices i. So henceforth we assume $0 < \alpha_i < 1$ and in particular the fact that $\alpha_i \geq 2^{-l}$.

Let $\delta = \frac{\rho - 1}{2\rho}$. Note $\delta > 2^{-\mathrm{poly}(n,l)}$.

Let $\epsilon = 2^{-(l+3)} w_{\min} \left(\frac{\rho-1}{\rho}\right)^2$. We will set $z_i = 1 - \delta w_i$ and first show that $z_i - \alpha_i - (1-\alpha_i)z_i \sum_j M_{ij}z_j$ is at least 2ϵ. Consider

$$\begin{aligned}
& z_i - \alpha_i - (1-\alpha_i)z_i \sum_j M_{ij} z_j \\
&= 1 - \delta w_i - \alpha_i - (1-\alpha_i)(1-\delta w_i)\sum_j M_{ij}(1 - \delta w_j) \\
&= 1 - \delta w_i - \alpha_i - (1-\alpha_i)(1-\delta w_i)(1 - \delta \sum_j M_{ij} w_j) \\
&= (1-\delta w_i)\left(\delta \sum_j (1-\alpha_i) M_{ij} w_j\right) - \delta w_i \alpha_i \\
&\geq (1-\delta w_i)(\delta \rho \alpha_i w_i) - \delta w_i \alpha_i \\
&= \delta \alpha_i w_i (\rho - \rho \delta w_i - 1) \\
&\geq \delta \alpha_i w_i \rho \\
&\geq \left(\frac{\rho - 1}{2\rho}\right)^2 \alpha_i w_i \\
&\geq 2\epsilon.
\end{aligned}$$

Now consider any vector $\vec{y}$ such that $z_i - 2\epsilon \leq y_i \leq z_i$. We claim that $\vec{y}$ is feasible. First, $y_i \leq 1$ since $y_i \leq z_i = 1 - \delta w_i \leq 1$. We now show that $y_i \geq 0$. First, $z_i \geq 0$ since $w_i \leq 1$ and $\delta < 1$. Since, as we showed above, $z_i - \alpha_i - (1-\alpha_i)z_i \sum_j M_{ij}z_j \geq 2\epsilon$, it follows that $y_i \geq z_i - 2\epsilon \geq \alpha_i + (1-\alpha_i)z_i \sum_j M_{ij}z_j \geq 0$. Finally,

$$\begin{aligned}
& y_i - \alpha_i - (1-\alpha_i)y_i \sum_j M_{ij} y_j \\
&\geq z_i - 2\epsilon - \alpha_i - (1-\alpha_i)y_i \sum_j M_{ij} y_j \\
&\geq z_i - 2\epsilon - \alpha_i - (1-\alpha_i)z_i \sum_j M_{ij} z_j \\
&\geq 0 \quad \text{(Using the claim about the } z_i\text{'s.)}
\end{aligned}$$

Thus setting $x_i = z_i - \epsilon$, we note that every vector $\vec{y}$ satisfying $x_i - \epsilon \leq y_i \leq x_i + \epsilon$ is feasible. This concludes the proof. ■

Proof:[of Theorem 4.4] Given M, $\vec{\alpha}$ and ϵ, let β be as given by Lemma 4.19. We first compute a β-approximation to the vector of revocation probabilities in time $\mathrm{poly}(n, l, \log \frac{1}{\beta}) = \mathrm{poly}(n, l, \log \frac{1}{\epsilon})$ using Lemma 4.21. The output is a vector $\vec{r}'$ of $l' = \mathrm{poly}(n, l, \log \frac{1}{\epsilon})$-bit rationals. Applying Lemma 4.19 to M, $\vec{\alpha}$, $\vec{r}$ and ϵ, we obtain an ϵ-approximation to the limit probability vector π in time $\mathrm{poly}(n, l, l', \log \frac{1}{\epsilon}) = \mathrm{poly}(n, l, \log \frac{1}{\epsilon})$. ■

References

[1] R. Barrett, P.P. Maglio, and D.C. Kellem. How to personalize the web. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97).*

[2] M. Charikar, S.R. Kumar, P. Raghavan, S. Rajagopalan and A. Tomkins. On targeting Markov segments. *Proceedings of the ACM Symposium on Theory of Computing, 1999.*

[3] M. Grotschel, L. Lovasz, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimizations.* Springer-Verlag, Berlin, 1987.

[4] R.A. Horn and C.A. Johnson. *Matrix Analysis*, Cambridge University Press, 1985.

[5] S. Karlin and H. Taylor *A First Course in Stochastic Processes*, Academic Press, 1975.

[6] J.G. Kemeny, J.L. Snell, and A.W. Knapp. *Denumerable Markov Chains*. The University Series in Higher Mathematics. Van Nostrand, Princeton, NJ, 1966.

[7] H. Lieberman. An agent that assists web browsing. *Proceedings of IJCAI*, 1995.

[8] H. Minc. *Nonnegative Matrices*. John Wiley and Sons, New York, 1988.

[9] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, 1994.

[10] T. Selker. COACH: a teaching agent that learns. *Communications of the ACM* 37:7, 1994.