Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

BENJAMIN HILLER   TJARK VREDEVELD

Stochastic dominance analysis of Online Bin Coloring algorithms

ZIB-Report 12-42 (November 2012)

Herausgegeben vom
Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7
D-14195 Berlin-Dahlem

Telefon: 030-84185-0
Telefax: 030-84185-125

e-mail: bibliothek@zib.de
URL: http://www.zib.de

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Stochastic dominance analysis of Online Bin Coloring algorithms*

Benjamin Hiller
Zuse Institute Berlin
Takustraße 7
D–14195 Berlin, Germany.
hiller@zib.de

Tjark Vredeveld
Maastricht University
Faculty of Economics and Business Administration
6200 MD Maastricht, The Netherlands.
t.vredeveld@maastrichtuniversity.nl

November 21, 2012

Abstract

This paper proposes a new method for probabilistic analysis of online algorithms. It is based on the notion of stochastic dominance. We develop the method for the online bin coloring problem introduced in [18]. Using methods for the stochastic comparison of Markov chains we establish the result that the performance of the online algorithm GREEDYFIT is stochastically better than the performance of the algorithm ONEBIN for any number of items processed. This result gives a more realistic picture than competitive analysis and explains the behavior observed in simulations.

1 Introduction

In *online optimization problems*, the input data is revealed step-by-step and an *online algorithm* has to take decisions that cannot be revoked once more data becomes known. In contrast, in classical or *offline optimization problems*, all input data is available a priori and therefore an offline algorithm can find better solutions than an online one.

Competitive analysis, introduced by [27] and named by [15], is by now the standard yardstick to evaluate online algorithms. It measures the quality of an online algorithm by comparing its solution value to that of an optimal offline algorithm that has full knowledge of the future. Due to its worst-case nature, competitive analysis sometimes fails to discriminate between algorithms or gives an overly pessimistic

*Supported by the DFG research group "Algorithms, Structure, Randomness" (Grant number GR 883/10-3, GR 883/10-4) and a DAAD dissertation grant. An extended abstract with parts of this work appeared in [12].

view. Therefore there is a need for alternatives to standard competitive analysis. Various alternatives have been proposed in the literature; almost all of them are based on the idea of weakening the offline adversary. This can be done by considering randomized online algorithms [6] or by allowing the online algorithms to use more resources [14]. Approaches which consider randomized input sequences include the diffuse adversary [16], average-case competitive analysis [25] and smoothed competitive analysis [5]. Direct comparisons of online algorithms are done in comparative analysis [16], the relative worst order ratio [7], and bijective analysis [1]. For a recent survey on probabilistic alternatives, see [13].

Stochastic dominance analysis. In this paper, we propose a new method for probabilistic analysis of online algorithms. In our approach we compare two algorithms directly, i. e., without reference to a hypothetical optimal offline algorithm. For this comparison we use the concept of stochastic dominance. Stochastic comparison methods have been successfully applied in areas like queueing theory [26], finance, economics and in particular decision under risk [23]. In this paper we introduce these concepts to the study of online algorithms.

We analyse the performance of online algorithms on random input sequences. The value of a solution obtained by an online algorithm is then described by a random variable. To compare two algorithms, one can compare the induced random variables. The most common way to do this is to compare their expected values, which results in *average-case analysis*. In a sense, this is the simplest form of a stochastic order. A more distinctive approach is to use a stronger stochastic order: the stochastic dominance order. A random variable X is *stochastically dominated* by a random variable Y, denoted by $X \leq_{\mathrm{st}} Y$, if

$$\Pr[X \geq x] \leq \Pr[Y \geq x] \quad \text{for all } x \in \mathbb{R}.$$

To apply this to the analysis of online algorithms, suppose the solution values of two online algorithms $\mathscr{A}$ and $\mathscr{B}$ are given by random variables $\chi^{\mathscr{A}}$ and $\chi^{\mathscr{B}}$, respectively. We then say that $\mathscr{A}$ is *stochastically better* than $\mathscr{B}$ (for a minimization problem), if $\chi^{\mathscr{A}} \leq_{\mathrm{st}} \chi^{\mathscr{B}}$.

Notice that using stochastic dominance analysis, we compare probability distribution functions instead of aggregated statistics like the maximum or the average. In this respect it is similar to bijective analysis [1], which is in fact a special case of stochastic dominance analysis.

We remark that stochastic dominance has been used for the analysis of algorithms before, see e. g., [22]. So far it was used to simplify parts of an analysis, e. g., to establish that a certain easier-to-analyze model provides bounds on the expectation of the real model studied. However, our usage is different: We use stochastic dominance to assess the relative quality of two online algorithms.

Our contribution. We propose a new probabilistic analysis of both deterministic and randomized online algorithms. As far as we know, this is the first use of stochastic dominance in the analysis of the quality of online and approximation algorithms.

We apply this new probabilistic analysis on online algorithms for the so-called *bin coloring* problem introduced by Krumke et al. [18], in which unit sized items of different colors need to be packed in bins of capacity B so as to minimize the maximum number of distinct colors in a bin. We will refer to the number of different colors in a bin as its *colorfulness*. Krumke et al. showed the counter-intuitive result that the trivial algorithm that packs the items in one bin until the bin is full outperforms a greedy algorithm in terms of competitive ratio. However, simulations show that the greedy algorithm actually gives significantly better results than the trivial one.

In our analysis, we assume that the color sequence is generated by choosing the colors i. i. d. according to a color distribution γ. In this model, all online algorithms eventually have to produce a bin with colorfulness equal to the capacity of the bins when the number of colors is sufficiently high. This implies that in this case, all online algorithms are asymptotically equally bad. Moreover, the asymptotic competitive ratio is 1 with probability 1. Both issues indicate that asymptotic probabilistic analysis does not give meaningful results.

Therefore, we compare the *transient* quality of two online algorithms for the bin coloring problem using stochastic dominance. We can show that the greedy algorithm is stochastically better than the trivial algorithm on any length of the sequence and any color distribution γ. We also obtain a similar result when the objective is the sum of colorfulnesses of the bins instead of the maximum colorfulness. Both results are based on an analysis of Markov chains related to the algorithms.

We emphasize several implications of this result. The fact that online algorithm $\mathscr{A}$ is stochastically better than online algorithm $\mathscr{B}$ implies not only that the expected value of $\mathscr{A}$ is bounded by that of $\mathscr{B}$, but also that the average competitive ratio of $\mathscr{A}$ is not more than that of $\mathscr{B}$. The average competitive ratio is defined as the ratio of the expected value of the online algorithm to that of the offline optimum. Furthermore it implies that $\mathscr{A}$ remains superior to $\mathscr{B}$ for all other performance measures that are nondecreasing functions of the considered objective function, which is of particular interest if the online optimization problem is actually a subproblem in some bigger context. By considering the uniform color distribution, our result can be interpreted as a counting result stating that there are more instances for which $\mathscr{A}$ manages to achieve a low maximum colorfulness than for $\mathscr{B}$.

Structure of the Paper. Section 2 defines the problem variants and discusses an application to elevator control. In Section 3 we present our new approach for the analysis of online algorithm based on stochastic dominance. We apply this approach to the analysis of bin coloring algorithms in Section 4. Section 5 discusses briefly some additional computer proof results and concludes the paper. In the appendix we present details on why existing proof techniques for stochastic dominance are not applicable here and how one of these techniques can be extended to obtain the computer proofs.

2 Problem definition and an application to elevator control

An instance of the *bin coloring problem* is described by the number of simultaneously open bins m, the bin capacity B and a sequence of n unit sized items, each of which has one of C colors. The items need to be packed in the open bins and whenever a bin has B items, it is closed and replaced by a new empty bin. The colorfulness of a bin is the number of different colors in this bin. We consider two objective functions: The original one considered in [18] is to minimize the maximum colorfulness of all bins. In addition, we also consider minimizing the sum of colorfulnesses of the bins, which is equivalent to minimizing the average colorfulness. We will refer to these variants as max-BC and sum-BC, respectively. In the online versions of both variants, the items arrive one by one and must be irrevocably assigned to a bin before the next item becomes known.

The bin coloring problem was originally motivated by an application to automated order picking [18]. Since then it has found applications in vehicle routing [24], networking [19], and elevator control, as described at the end of this section. It is shown in [24, 19, 18] that the bin coloring problem is NP-hard and that it cannot be approximated in polynomial time within a factor of $4/3$ unless $\mathrm{P} = \mathrm{NP}$. In [19] and [18] algorithms are provided that find a solution of cost $\mathrm{OPT} + 1$ in the case that there are exactly mB items.

We consider the following bin coloring algorithms.

GREEDYFIT packs an item having a color that is already in one of the open bins, in the bin carrying that color and, if the color is in none of the open bins, then chooses a bin which currently has the least number of different colors.

ONEBIN uses only a single bin, i. e., it packs all items in the same bin until it is full and replaced.

FIXEDCOLORS uses a prescribed color-to-bin-assignment to put the items into bins according to their colors.

Observe that FIXEDCOLORS is somewhere between ONEBIN and GREEDYFIT: it takes advantage of all available bins, but uses a static assignment rule of colors to the bins. The main part focuses on results about ONEBIN and GREEDYFIT; results about FIXEDCOLORS can be found in the appendix.

Krumke et al. [18] showed that the competitive ratio of ONEBIN is at most $2m-1$ and that of GREEDYFIT is at least $2m$. That is, they established the counterintuitive result that in terms of the competitive ratio, the trivial algorithm ONEBIN is better than the more sophisticated algorithm GREEDYFIT. The authors mentioned that the most challenging issue is to analyze the algorithms from an average-case point of view to explain the clear dominance of GREEDYFIT over ONEBIN observed in simulations. Such an average-case analysis is a consequence of our results.

For probabilistic analysis, we assume that the number of open bins m, the bin capacity B and the number of colors C is given deterministically. The color sequence, however, is generated by chosing each color independently according to a probability distribution function γ over the colors.

Note that in this model, all online algorithms eventually have to produce a bin with colorfulness B if the number of colors is sufficiently high, say $C \geq 2mB$. This implies that in this case, all online algorithms are asymptotically equally bad. Moreover, since eventually there will be a color subsequence of length $2mB$ with all colors different, the asymptotic competitive ratio is 1 with probability 1. Both issues indicate that asymptotic probabilistic analysis does not give meaningful results. Instead it is necessary to study the *transient* behavior instead of the asymptotic behavior of the algorithms.

Let the random variables χ_n^{GF} and χ_n^{OB} denote the maximum colorfulness attained after processing n items using GREEDYFIT and ONEBIN, respectively. We will show that GREEDYFIT is stochastically better than ONEBIN after n items, i. e., that $\chi_n^{\mathrm{GF}} \leq_{\mathrm{st}} \chi_n^{\mathrm{OB}}$. The competitive analysis results of Krumke et al. [18] imply that there is a sequence on which ONEBIN gives better overall colorfulness than GreedyFit, i. e., GREEDYFIT is not sequence-wise at least as good as ONEBIN.

Application to elevator control. Before starting our stochastic dominance analysis, we briefly explain an interesting connection between the online bin coloring problem and the relative performance of conventional 2-button elevator control and elevator control based on *destination calls*. In the latter system, a passenger registers his or her destination floor right at his start floor instead of the travel direction only. This way, the elevator control has more information as a basis for its decisions which hopefully leads to better performance.

Let us consider a group of L elevators operating in an office building with N floors. Elevator engineers consider different traffic patterns for analyzing elevator control performance. The most import one is *up peak traffic*: In the morning, there is a high-intensity traffic from the entrance floors to the office floors. Up peak traffic is considered to be the most demanding traffic situation [4], so it is often used for dimensioning an elevator system. An important figure in this context is the (up peak) handling capacity HC [4]. The handling capacity is defined as the critical passenger arrival rate: If the passenger arrival rate is higher than the handling capacity, the system cannot cope with the resulting traffic and waiting times increase rapidly. The handling capacity is inversely related to the time needed for a single roundtrip, RTT. Elevator engineers estimate RTT using the formula

$$RTT = 2Ht_v + (S+1)t_s + Pt_l,$$

where H is the highest floor reached, t_v the time needed to pass one floor, S is the number of stops needed to unload the passengers, t_s the time per stop, P is the cabin capacity, and t_l the time to unload one passenger. Note that only H and S depend on the actual load of the lift. For calculations, H is usually assumed to be N or $N-1$, which is justified when the cabin capacity P is high. The number of stops S is calculated by assuming some distribution of the destination floors.

Let us assume that the passenger arrival rate is so high that the elevators leaving the main entrance floor are always full (this is also assumed in the formula for the handling capacity above). Since we are only interested in the number of stops S for a roundtrip, we can model the passenger sequence by a sequence of destination floors only. Viewing the destination floors as colors and the elevators as bins, the situation can be seen as a bin coloring problem. The number of distinct colors is the number of floors in the building, N, whereas the bin capacity is equal to the capacity of an elevator, P. As H is usually assumed to be either N or $N-1$, to minimize RTT we need to minimize S, the number of stops of one elevator. That is, in terms of the bin coloring problem, we want to minimize the average colorfulness of the bins, thereby motivating the second objective that we consider.

We now want to compare the RTT for a conventional system to that of a destination call system using our bin coloring analysis. In a conventional system, passengers board the lifts in the order of their arrival at the main entrance floor. The stops of the resulting roundtrips are determined by the sub-sequences of size P of the destination floor sequence. That is, the operating of a conventional system can be modeled by ONEBIN.

For a destination call system the elevator control has the possibility to reduce stops by assigning passengers with the same destination floor to the same lift and balancing the number of stops between the L lifts. A natural strategy to do this is the GREEDYFIT algorithm, using up to L bins.

As we will see in Theorem 4.9, the average number of colors achieved by GREEDYFIT is stochastically dominated by that of ONEBIN. Applied to our elevator setting this means that the total number of stops in the destination call system is stochastically dominated by that of the conventional system, which implies the same relation for the RTT of both systems. Since $X \leq_{\text{st}} Y$ implies $1/X \geq_{\text{st}} 1/Y$, we get that the handling capacity of a destination call system is higher than that of a conventional system, independent of the destination floor distribution. Note that we could not conclude that the expected handling capacity is larger if we had shown only that the expected RTT is smaller.

3 Stochastic dominance and Markov chain models for online algorithms

Almost all probabilistic analyses of online algorithms up to now are more or less average case analyses, i. e., they focus on the *expected* performance. Sometimes results on the expected performance are used to obtain high probability or asymptotically almost surely results. Algorithms are compared by considering their expected performance, e. g., their expected competitive ratio. In this paper we propose to compare algorithms by looking at their *objective value distributions* using *stochastic orders*, in particular the stochastic dominance order. A random variable X is said to be *stochastically dominated* by a random variable Y, written $X \leq_{\text{st}} Y$, if

$$\Pr[X \geq x] \leq \Pr[Y \geq x] \quad \text{for all } x \in \mathbb{R}. \tag{1}$$

Suppose we can describe the objective values of two online algorithms $\mathscr{A}$ and $\mathscr{B}$ after n steps by random variables $\chi_n^{\mathscr{A}}$ and $\chi_n^{\mathscr{B}}$, respectively. We then say that $\mathscr{A}$ is *stochastically better* than $\mathscr{B}$ w. r. t. to the considered input distribution if $\chi_n^{\mathscr{A}} \leq_{\text{st}} \chi_n^{\mathscr{B}}$ for all $n \in \mathbb{N}$ (assuming a minimization problem). Having a stochastic dominance result for two online algorithms has some nice consequences due to the properties of this stochastic order.

Theorem 3.1 (see e. g., [23]) *Let X and Y be two random variables with $X \leq_{\text{st}} Y$ and $f\colon \mathbb{R} \to \mathbb{R}$ be some nondecreasing function. Then*

- $\mathbb{E}[X] \leq \mathbb{E}[Y]$ *and*
- $f(X) \leq_{\text{st}} f(Y)$.

Suppose $(X_n)_{n\in\mathbb{N}_0}$ and $(Y_n)_{n\in\mathbb{N}_0}$ are sequences of random variables that converge in distribution to X and Y, respectively, and that satisfy $X_n \leq_{\text{st}} Y_n$. Then

- $X \leq_{\text{st}} Y$.

Assuming that the random variables X and Y measure the performance of algorithms $\mathscr{A}$ and $\mathscr{B}$, respectively, the first part of Theorem 3.1 states that if algorithm $\mathscr{A}$ is stochastically better than $\mathscr{B}$ it is also better in expectation. It also implies that $\mathscr{A}$ has a better expected competitive ratio than $\mathscr{B}$, where the *expected competitive ratio* of an online algorithm ALG is defined as the smallest c such that

$$\mathbb{E}[\text{ALG}] \leq c \cdot \mathbb{E}[\text{OPT}]. \tag{2}$$

The second part of the theorem states that $\mathscr{A}$ remains superior to $\mathscr{B}$ for all other performance measures that are nondecreasing functions of the original performance measure. Finally, the third part says that if $\mathscr{A}$ is stochastically better than $\mathscr{B}$ for any point in time, then it is also better in the limit.

Stochastic dominance results admit an interesting interpretation if the input sequences are uniformly distributed. It is then possible to view the probabilistic result deterministically as a counting result. Stochastic dominance then implies that the better algorithm achieves a low cost on more sequences than the worse one, i. e., it is more robust w. r. t. the unknown future, a property that is certainly desirable for an online algorithm. In fact, stochastic dominance in the special case of the uniform distribution is equivalent to a recent deterministic way to compare online algorithms, known as *bijective analysis* [1, 2]. Let S_n denote the input sequences of length n and consider two online algorithms ALG_1 and ALG_2. ALG_1 is said to dominate ALG_2 w. r. t. bijective analysis, if there is a bijective mapping $\phi\colon S_n \to S_n$ such that $\mathrm{ALG}_1(\sigma) \leq \mathrm{ALG}_2(\phi(\sigma))$ for any $\sigma \in S_n$. Intuitively, the input sequences are "remapped" such that ALG_1 is "pointwise" better than ALG_2.

It is quite natural to model the working of an online algorithm on random input as a (time homogeneous, discrete time) Markov chain, where the state space is given by the configurations an online algorithm generates while processing the input sequence. Formally, we will consider the following Markov chain model for online algorithms for minimization problems. A *valued Markov chain* (X, χ) is a Markov chain $X = (X_n)_{n\in\mathbb{N}_0}$ on state space $\mathscr{S}$ together with a valuation function $\chi\colon \mathscr{S} \to V$ for some $V \subseteq \mathbb{N}_0$. The random successor state of state $s \in \mathscr{S}$ is denoted by $X(s)$.

We now present valued Markov chain models for max-BC and sum-BC algorithms that are used for our stochastic dominance results in the next section. The state-space of the valued Markov chains is determined by the instance, i. e., the parameters m, B, and C, whereas the transitions model the working of the algorithm and are thus algorithm-dependent. Thus depending on the algorithm only a subset of the full state-space may be reached. The probabilities of the transitions arise due to the color distribution γ.

We will use the following valued Markov chain model for max-BC algorithms. The operation of any such bin coloring algorithm can be described on a state space which encodes for every bin i its current number of items f_i and the set of colors in that bin, C_i. Moreover, the state also keeps track of the maximum colorfulness attained so far. Formally, we have

$$\mathscr{S}_{\text{max-BC}} := \mathscr{S}_{\text{max-BC}}(m,B,C) = \Big\{(f_1,C_1,\ldots,f_m,C_m,\chi) \;\Big|\; 0 \leq |C_i| \leq f_i < B,$$
$$|C_i| \leq \chi \leq \min\{B,C\}\Big\}.$$

We will use $f_i(s)$, $C_i(s)$, and $\chi(s)$ to refer to the components of state s. Additionally, we define $c_i(s) := |C_i(s)|$. The state $(0,\emptyset,\ldots,0,\emptyset,0)$ is called the *initial empty state*.

Suppose an online bin coloring algorithm $\mathscr{A}$ is in state s and receives an item of color c. The algorithm then changes to state s' by putting this item in one of the bins, say bin i. There are two cases: Either color c is contained in $C_i(s)$, we say the color is *known (in bin i)*, or it is not, so the color is *new (in bin i)*. We will denote the successor state for the first case by $s^{\mathrm{k}(i)}$ (the color c is not needed to determine the successor state), for the second by $s^{\mathrm{n}(i,c)}$. It will be convenient not to consider the new color c, but to deal with the random state resulting from s if any new color distributed according to γ is seen. We will use the notation $s^{\mathbf{n}(i)}$ for this random state.

The OneBin algorithm is then described by the transitions

$$s' = \begin{cases} s^{\mathrm{k}(1)} & \text{with probability } \gamma(C_1(s)), \\ s^{\mathbf{n}(1)} & \text{with probability } 1-\gamma(C_1(s)), \end{cases} \tag{3}$$

where we use the shortcut notation $\gamma(S) := \sum_{s \in S} \gamma(s)$. This defines a Markov chain which we denote by $\mathrm{OB}(m, B, C, \gamma)$. Note that although ONEBIN uses only the first bin, we consider $\mathrm{OB}(m, B, C, \gamma)$ as working on the whole state space with m bins.

Similarly, we can give a Markov chain $\mathrm{GF}(m, B, C, \gamma)$ for GREEDYFIT. $\mathrm{GF}(s)$ is the bin GREEDYFIT selects for an item with a new color in state s. Depending on the tie-breaking rule used by the specific variant of GREEDYFIT, $\mathrm{GF}(s)$ may or may not be a random variable. We only need that $\mathrm{GF}(s)$ is one of the bins having in state s the smallest number of colors. The successor state s' of s is given by

$$s' = \begin{cases} s^{\mathrm{k}(i)} & \text{with probability } \gamma(C_i(s)) \qquad 1 \le i \le m, \\ s^{\mathbf{n}(\mathrm{GF}(s))} & \text{with probability } 1 - \gamma\big(\bigcup_i C_i(s)\big). \end{cases} \tag{4}$$

The operation of online algorithms in the sum-BC problem can be captured by a similar Markov chain model. The main difference is that the χ-component is no longer the maximum of the colorfulness seen so far, but the sum. Note that the resulting Markov chains are infinite. We can use the state space

$$\mathscr{S}_{\text{sum-BC}} := \mathscr{S}_{\text{sum-BC}}(m, B, C) = \Big\{ (f_1, C_1, \dots, f_m, C_m, \chi) \;\Big|\; 0 \le |C_i| \le f_i < B, \\ |C_i| \le \min\{B, C\}, \chi \in \mathbb{N}_0 \Big\}.$$

The χ-component increases each time a new color for a bin is encountered.

To avoid notational overhead, we will use the same notation for both problem variants. Therefore the sum-BC-Markov chains for ONEBIN and GREEDYFIT will be denoted by $\mathrm{OB}(m, B, C, \gamma)$ and $\mathrm{GF}(m, B, C, \gamma)$, too. We use the notations $\mathrm{OB}(m, B, C, \gamma)_n$ and $\mathrm{GF}(m, B, C, \gamma)_n$ for the random state after n steps when ONEBIN and GREEDYFIT are started in the initial empty state.

In our setting there are basically two ways to measure the performance of an online algorithm. The first, more natural one, was already introduced: We look at the evolution of the valuation over time, i. e., at $\chi(X_n)$ if (X, χ) is a valued Markov chain modelling the algorithm. But we may also ask: How fast does the valuation grow, i. e., for how many steps can the algorithm "guarantee" that the valuation is at most $v \in V$? Of course, an algorithm is regarded to be good if it keeps a low valuation for a long time. Formally, we look at the stopping time T_X^v, a random variable that gives the first time the Markov chain X reaches a state with valuation at least v. The following result states that both performance measures are actually equivalent if we compare algorithms with the stochastic dominance order. The theorem is an extension of a result in [21].

Theorem 3.2 *Let (X, ϕ) and (Y, ψ) be valued Markov chains on state spaces $\mathscr{S}_X$ and $\mathscr{S}_Y$ with common valuation space $V \subseteq \mathbb{N}_0$. Assume that the transitions of X and Y are such that the value of a state is nondecreasing in each step and that $\phi(X_0) = \psi(Y_0)$. Then the following are equivalent:*

1. $T_Y^v \le_{\mathrm{st}} T_X^v \quad \forall v \in V$.
2. $\phi(X_n) \le_{\mathrm{st}} \psi(Y_n) \quad \forall n \in \mathbb{N}_0$.

Proof. Let the Markov chain X be defined on the probability space $(\Omega, \mathscr{A}, \Pr)$. The stopping time T_X^v is then a random variable $T_X^v \colon \Omega \to \mathbb{N}_0$ that is defined by

$$T_X^v(\omega) := \min \big\{ n \,\big|\, \phi\big(X_n(\omega)\big) \ge v \big\}$$

for each $\omega \in \Omega$. Since $\phi(X_n(\omega)) \geq \phi(X_{n'}(\omega))$ whenever $n \geq n'$, we have the equivalence

$$T_X^v(\omega) \leq n \iff \phi(X_n(\omega)) \geq v,$$

which implies

$$\Pr[T_X^v \leq n] = \Pr[\phi(X_n) \geq v].$$

Of course, analogous statements hold for Y as well.

We now have the following chain of equivalences.

$$\begin{aligned} &\phi(X_n) \leq_{\mathrm{st}} \psi(Y_n) \quad \forall n \in \mathbb{N}_0 \\ \iff &\Pr[\phi(X_n) \geq v] \leq \Pr[\psi(Y_n) \geq v] \quad \forall n \in \mathbb{N}_0, v \in V \\ \iff &\Pr[T_X^v \leq n] \leq \Pr[T_Y^v \leq n] \quad \forall n \in \mathbb{N}_0, v \in V \\ \iff &T_Y^v \leq_{\mathrm{st}} T_X^v \quad \forall v \in V. \end{aligned}$$

4 Stochastic dominance between ONEBIN and GREEDYFIT

The goal of this section is to show that, in both problem variants,

$$\chi\big(\mathrm{GF}(m,B,C,\gamma)_n\big) \leq_{\mathrm{st}} \chi\big(\mathrm{OB}(m,B,C,\gamma)_n\big) \quad \forall n$$

for all parameters m,B,C and color distributions γ. This kind of result is known as *comparison result* for Markov chains in the probability theory literature, see e. g., [23].

Unfortunately, the general comparison results for Markov chains based on stochastic monotonicity [23] are not sufficient to prove stochastic dominance between GREEDYFIT and ONEBIN. Doisy [9] developed a comparison criterion that is not based on stochastic monotonicity, however, this result is also too weak. A detailed discussion can be found in the appendix.

4.1 Preliminaries

Consider valued Markov chains (X,χ) and (Y,χ) corresponding to bin coloring algorithms, where χ measures either the maximum or the sum of the colorfulnesses encountered. Recall that for a valued Markov chain (X,χ), we denote by T_X^v the first time the Markov chain X reaches a state with valuation at least v. By Theorem 3.2, it is suffient to show $T_Y^v \leq_{\mathrm{st}} T_X^v$ for all $v \in V$ to show $\chi(X_n) \leq_{\mathrm{st}} \chi(Y_n)$ for all $n \in \mathbb{N}_0$. We will therefore analyze the stopping times to prove our stochastic dominance result. In the sequel, we denote by $T_X^v(s)$ the stopping time for reaching a state with valuation at least v when started deterministically in state in s.

How can we show $T_Y^v(s_0) \leq_{\mathrm{st}} T_X^v(s_0)$? In order to apply a kind of induction technique we introduce a family of Markov chains $\big(X(n)\big)_{n\in\mathbb{N}}$ derived from a Markov chain X as follows. The state space of $X(n)$ is $\mathscr{S} \times \{0,\dots,n\}$ and the transitions are defined by

$$\begin{aligned} \Pr\big[X(n)_{i+1} = (s',i+1) \mid X(n)_i = (s,i)\big] &:= \Pr\big[X_{i+1} = s' \mid X_i = s\big] && \forall 0 \leq i < n, \\ \Pr\big[X(n)_{i+1} = (s,n) \mid X(n)_i = (s,n)\big] &:= 1 && \forall i \geq n. \end{aligned}$$

The Markov chain $X(n)$ can be thought of as an time-expanded, acyclic version of the chain X for the first n steps. Clearly, we have

$$\Pr\left[T_X^v(s) = i\right] = \Pr\left[T_{X(n)}^v((s,0)) = i\right] \quad \forall 0 \le i < n. \tag{5}$$

So in order to show $T_Y^v(s_0) \le_{\mathrm{st}} T_X^v(s_0)$, we can prove that

$$T_{Y(n)}^v((s_0,0)) \le_{\mathrm{st}} T_{X(n)}^v((s_0,0)) \quad \forall n \in \mathbb{N}.$$

To simplify notation, we will write $T_{X(n)}^v(s)$ for $T_{X(n)}^v((s,0))$ from now on. We have the following simple result.

Lemma 4.1 *For any valued Markov chain (X,χ) the stochastic dominance relation*

$$T_{X(n+1)}^v(s) \le_{\mathrm{st}} T_{X(n)}^v(s)$$

holds for all states s, $n \in \mathbb{N}_0$, and $v \in V$.

Proof. Consider a sample path $\omega = (X_0, X_1, \ldots, X_n)$ of $X(n)$ with $X_0 = s$. Obviously, ω can be extended to a sample path $\omega' = (X_0, X_1, \ldots, X_n, X_{n+1})$ of $X(n+1)$ and all sample paths of $X(n+1)$ starting in s are obtained this way. There are two cases:

1. $\chi(X_n) \ge v$: For all ω' that are extensions of ω we have $T_{X(n+1)}^v(s) = T_{X(n)}^v(s)$.
2. $\chi(X_n) < v$: For those ω, $T_{X(n)}^v(s)$ is infinite, whereas for any extension ω' $T_{X(n+1)}^v(s)$ is either $n+1$ or infinite, too. Thus $T_{X(n+1)}^v(s) \le T_{X(n)}^v(s)$.

To analyze the stopping times, we will employ the concept of a mixture of random variables.

Definition 4.2 Let $(X_m)_{m\in M}$ be a family of random variables and Θ be an M-valued random variable. The random variable Y defined by $Y := X_\Theta$, i. e., the X-variable to use is given by the realization of Θ, is called a *mixture* and denoted by $[(X_m)_{m\in M} \mid \Theta]$.

An important property of $\le_{\mathrm{st}}$ is that it is closed under mixture, as stated in the Mixture Theorem.

Theorem 4.3 ([23, p. 6]) *Suppose $[(X_m)_{m\in M} \mid \Theta]$ and $[(Y_m)_{m\in M} \mid \Theta]$ are two mixtures controlled by the same random variable Θ satisfying $X_m \le_{\mathrm{st}} Y_m$ for all $m \in M$. Then we have*

$$[(X_m)_{m\in M} \mid \Theta] \le_{\mathrm{st}} [(Y_m)_{m\in M} \mid \Theta].$$

For two random variables X and Y, we will frequently write $X = Y$ to mean that they have the same distribution function.

4.2 GreedyFit is better than OneBin: max-BC

We will now apply the strategy described above to the comparison of the GreedyFit and the OneBin bin coloring algorithms. The main technique is to analyze a kind of stochastic recursion for $T_{X(n)}^v$ based on a mixture of random variables.

Let $\mathrm{OB} = \mathrm{OB}(m,B,C,\gamma)$ for fixed parameters m,B,C,γ. In a state $s \in \mathscr{S}_{\text{max-BC}}$ OneBin does the transitions to states

$$\begin{cases} s^{\mathrm{k}(1)} & \text{with probability } \gamma(C_1(s)), \\ s^{\mathbf{n}(1)} & \text{with probability } 1 - \gamma(C_1(s)). \end{cases}$$

Using the random variable $\Theta\colon \mathscr{S}_{\text{max-BC}} \to \mathbb{N}$ defined by

$$\Theta(s) := \begin{cases} 1 & \text{the next color is known in bin 1,} \\ 2 & \text{the next color is new in bin 1,} \end{cases}$$

we can come up with a recursive expression for $T^{v}_{\mathrm{OB}(n)}(s)$, namely

$$T^{v}_{\mathrm{OB}(n)}(s) = \begin{cases} 0 & \chi(s) \geq v, \\ 1 + \left[T^{v}_{\mathrm{OB}(n-1)}\big(s^{\mathrm{k}(1)}\big), T^{v}_{\mathrm{OB}(n-1)}\big(s^{\mathbf{n}(1)}\big) \,\middle|\, \Theta(s) \right] & \chi(s) < v. \end{cases} \tag{6}$$

This recursion and the Mixture Theorem 4.3 are the most important ingredients for the proofs to come.

We call two states $s, s' \in \mathscr{S}_{\text{max-BC}}$ OB-*equivalent*, if the valuation, the number of items and the set of colors in bin 1 are the same in s and s', i. e., if $\chi(s) = \chi(s')$, $f_1(s) = f_1(s')$, and $C_1(s) = C_1(s')$. Note that OneBin behaves exactly the same in two OB-equivalent states and therefore the stopping times from two OB-equivalent states coincide. The following lemma gives some useful comparisons of stopping times from certain states in the $\mathrm{OB}(n)$ chains.

Lemma 4.4 *Consider the* OneBin *Markov chain* $\mathrm{OB} = \mathrm{OB}(m,B,C,\gamma)$ *for parameters* $m,B \geq 2$*,* C*, and color distribution* γ*. We have for all states* $s \in \mathscr{S}_{\text{max-BC}}$*,* $n \in \mathbb{N}$*, and* $v \in V$*:*

1. $T^{v}_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}) \leq_{\mathrm{st}} T^{v}_{\mathrm{OB}(n)}(s^{\mathrm{k}(1)})$*, and*
2. $T^{v}_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}) \leq_{\mathrm{st}} T^{v}_{\mathrm{OB}(n)}(s')$ *for every state* s' *that is OB-equivalent to* s*.*

Proof. Let $s \in \mathscr{S}_{\text{max-BC}}$ be such that $\chi(s^{\mathbf{n}(1)}) < v$ (the case $\chi(s^{\mathbf{n}(1)}) \geq v$ is trivial).

1. Observe that both $s^{\mathbf{n}(1)}$ and $s^{\mathrm{k}(1)}$ have the same number of items in bin 1, say f. In the case $f = 0$ both states are OB-equivalent, since bin 1 is empty then. Therefore, $T^{v}_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}) = T^{v}_{\mathrm{OB}(n)}(s^{\mathrm{k}(1)})$.

   For $f > 0$ the evolution from both states will be identical after $B - f$ steps, since bin 1 is then empty again. It is therefore sufficient to show

   $$T^{v}_{\mathrm{OB}(B-f)}(s^{\mathbf{n}(1)}) \leq_{\mathrm{st}} T^{v}_{\mathrm{OB}(B-f)}(s^{\mathrm{k}(1)})$$

   for all $1 \leq f < B$. We will show this by induction on $j := B - f$.

   To start the induction, consider $j = 1$. There are two cases:

   - $c_1(s) = v - 2$: We have $\Pr\left[T^{v}_{\mathrm{OB}(j)}(s^{\mathbf{n}(1)}) \leq 1\right] \leq 1$ and $\Pr\left[T^{v}_{\mathrm{OB}(j)}(s^{\mathrm{k}(1)}) = \infty\right] = 1$.
   - $c_1(s) < v - 2$: Then $T^{v}_{\mathrm{OB}(j)}(s^{\mathbf{n}(1)}) = T^{v}_{\mathrm{OB}(j)}(s^{\mathrm{k}(1)}) \equiv \infty$.

   In both cases, the stochastic dominance is immediate.

   Let us now consider the induction step, i. e., $j > 1$. The key observation is that (since we need at least two items to close bin 1)

   $$f_1(s^{\mathbf{n}(1),\mathrm{k}(1)}) = f_1(s^{\mathrm{k}(1),\mathbf{n}(1)}) \quad \text{and}$$
   $$C_1(s^{\mathbf{n}(1),\mathrm{k}(1)}) = C_1(s^{\mathrm{k}(1),\mathbf{n}(1)}),$$

which means that both states are OB-equivalent. Using the Mixture Theorem 4.3, we can then estimate

$$\begin{aligned}
&T^v_{\mathrm{OB}(j)}(s^{\mathbf{n}(1)}) \\
&= 1 + \left[T^v_{\mathrm{OB}(j-1)}(s^{\mathbf{n}(1),\mathrm{k}(1)}), T^v_{\mathrm{OB}(j-1)}(s^{\mathbf{n}(1),\mathbf{n}(1)}) \,\middle|\, \Theta(s^{\mathbf{n}(1)})\right] \\
&\leq_{\mathrm{st}} 1 + T^v_{\mathrm{OB}(j-1)}(s^{\mathbf{n}(1),\mathrm{k}(1)}) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(by induction)} \\
&= 1 + T^v_{\mathrm{OB}(j-1)}(s^{\mathrm{k}(1),\mathbf{n}(1)}) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(OB-equivalence)} \\
&\leq_{\mathrm{st}} 1 + \left[T^v_{\mathrm{OB}(j-1)}(s^{\mathrm{k}(1),\mathrm{k}(1)}), T^v_{\mathrm{OB}(j-1)}(s^{\mathrm{k}(1),\mathbf{n}(1)}) \,\middle|\, \Theta(s^{\mathrm{k}(1)})\right] \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(by induction)} \\
&= T^v_{\mathrm{OB}(j)}(s^{\mathrm{k}(1)}).
\end{aligned}$$

2. Since s and s' are OB-equivalent, we have

$$\begin{aligned}
&T^v_{\mathrm{OB}(n)}(s') = T^v_{\mathrm{OB}(n)}(s) && \text{(OB-equivalence)} \\
&= 1 + \left[T^v_{\mathrm{OB}(n-1)}(s^{\mathbf{n}(1)}), T^v_{\mathrm{OB}(n-1)}(s^{\mathrm{k}(1)}) \,\middle|\, \Theta(s)\right] \\
&\geq_{\mathrm{st}} 1 + T^v_{\mathrm{OB}(n-1)}(s^{\mathbf{n}(1)}) && \text{(by 1.)} \\
&\geq_{\mathrm{st}} 1 + T^v_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}) && \text{(by Lemma 4.1)} \\
&\geq_{\mathrm{st}} T^v_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}).
\end{aligned}$$

Theorem 4.5 *Let* OB *and* GF *be the* ONEBIN *and* GREEDYFIT max-BC-*Markov chains for fixed parameters m, B, C with $B, m \geq 2$ for some color distribution γ. We have for all states $s \in \mathscr{S}_{\text{max-BC}}$, $n \in \mathbb{N}$, and $v \in V$:*

$$T^v_{\mathrm{OB}(n)}(s) \leq_{\mathrm{st}} T^v_{\mathrm{GF}(n)}(s).$$

Proof. The proof is by induction on n. Since GREEDYFIT is not worse than ONEBIN for a single step in each state s, we have $T^v_{\mathrm{OB}(1)}(s) \leq_{\mathrm{st}} T^v_{\mathrm{GF}(1)}(s)$.

The proof idea for the induction step is depicted in Figure 1. Suppose we know that $T^v_{\mathrm{OB}(n)}(s) \leq_{\mathrm{st}} T^v_{\mathrm{GF}(n)}(s)$ for all $s \in \mathscr{S}_{\text{max-BC}}$. Consider a state $s \in \mathscr{S}_{\text{max-BC}}$ having $\chi(s) < v$. Define the random variable $\Theta\colon \mathscr{S}_{\text{max-BC}} \to \{1, \ldots, m+1\}$ by

$$\Pr[\Theta(s) = i] = \begin{cases} \gamma(C_i(s)) & 1 \leq i \leq m, \\ 1 - \gamma\big(\bigcup_i C_i(s)\big) & i = m+1, \end{cases}$$

i. e., Θ in a sense "selects" the GREEDYFIT successor of state s. Using Θ, we can write the recursion for the stopping time of OB as

$$\begin{aligned}
&T^v_{\mathrm{OB}(n+1)}(s) \\
&= 1 + \left[T^v_{\mathrm{OB}(n)}(s^{\mathrm{k}(1)}), T^v_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}), \ldots, T^v_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}) \,\middle|\, \Theta(s)\right].
\end{aligned}$$

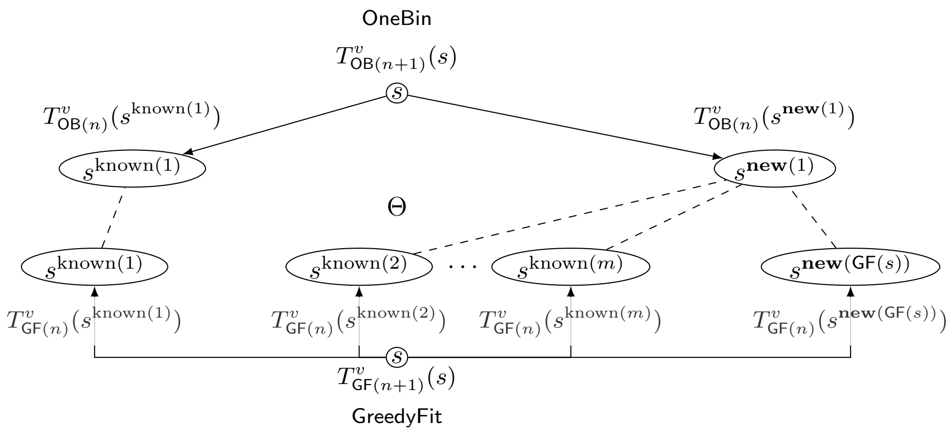

Figure 1: Proof idea for induction step in the proof of Theorem 4.5. The upper part shows the transitions of the OB Markov chain, the lower part those of the GF chain, both with corresponding stopping times. Using Lemma 4.4 and the induction hypothesis, one can map (dashed lines indicate the mapping Θ) the successor states of the OB chain to those of the GF chain such that we have stochastic dominance for each pair of stopping times.

Observe that $s^{\mathrm{k}(i)}$, $2 \le i \le m$, are OB-equivalent to s, $s^{\mathbf{n}(\mathrm{GF}(s))}$ is either OB-equivalent to s or equal to $s^{\mathbf{n}(1)}$. We use Lemma 4.4 to bound this by

$$\le_{\mathrm{st}} 1 + \left[T^v_{\mathrm{OB}(n)}(s^{\mathrm{k}(1)}), \ldots, T^v_{\mathrm{OB}(n)}(s^{\mathrm{k}(m)}), T^v_{\mathrm{OB}(n)}(s^{\mathbf{n}(\mathrm{GF}(s))}) \,\middle|\, \Theta(s)\right],$$

which by the induction hypothesis is bounded by

$$\begin{aligned}&\le_{\mathrm{st}} 1 + \left[T^v_{\mathrm{GF}(n)}(s^{\mathrm{k}(1)}), \ldots, T^v_{\mathrm{GF}(n)}(s^{\mathrm{k}(m)}), T^v_{\mathrm{GF}(n)}(s^{\mathbf{n}(\mathrm{GF}(s))}) \,\middle|\, \Theta(s)\right]\\ &= T^v_{\mathrm{GF}(n+1)}(s).\end{aligned}$$

This concludes the induction step and the proof.

Corollary 4.6 *Let* OB *and* GF *be the* ONEBIN *and* GREEDYFIT max-BC-*Markov chains for fixed parameters m, B, C and color distribution γ. We have for all states $s \in \mathscr{S}_{\text{max-BC}}$, in particular the initial empty state, and for all $n \in \mathbb{N}_0$ that*

$$\chi\big(\mathrm{GF}(s)_n\big) \le_{\mathrm{st}} \chi\big(\mathrm{OB}(s)_n\big).$$

Proof. The cases $m = 1$ and $B = 1$ are trivial. For the remaining cases, combine Theorems 3.2 and 4.5 and the relation of $\mathrm{OB}(n)$ and OB as well as $\mathrm{GF}(n)$ and GF according to Equation (5).

4.3 GREEDYFIT is better than ONEBIN: sum-BC

The analysis of the sum-BC problem is very similar to the one of max-BC in the preceding section. Recall that the state space of the sum-BC only differs from the one of the max-BC in its interpretation of the χ-component: it now counts the sum of the colorfulnesses of all used bins instead of the maximum. Therefore, the χ-component increases with every transition due to a new color. Nevertheless, recursion (6) for the stopping times is also valid for the analysis of the sum-BC.

Note that the proof of Theorem 4.5 is based only on Lemma 4.4. The proof of item 2 of Lemma 4.4 only needs item 1 and OB-equivalence. The notion of OB-equivalence introduced for the max-BC is also appropriate for the sum-BC. In particular, stopping times for two OB-equivalent states coincide also for the sum-BC-Markov chain of ONEBIN. Due to these observations, it is sufficient to prove an analogue of item 1 of Lemma 4.4 to establish stochastic dominance between GREEDYFIT and ONEBIN for the sum-BC. The proof uses the concept of a coupling Markov chain.

Definition 4.7 Let $X = (X_n)_{n\in\mathbb{N}_0}$ and $Y = (Y_n)_{n\in\mathbb{N}_0}$ be Markov chains on state spaces $\mathscr{S}_X$ and $\mathscr{S}_Y$, respectively. A Markov chain $Z = (\tilde{X}, \tilde{Y})$ on state space $\mathscr{S}_X \times \mathscr{S}_Y$ is a *coupling Markov chain* if $\tilde{X}$ and $\tilde{Y}$ are distributed as X and Y, respectively. However, $\tilde{X}$ and $\tilde{Y}$ need not be independent.

Lemma 4.8 *Consider the* ONEBIN *Markov chain* $\mathrm{OB} = \mathrm{OB}(m, B, C, \gamma)$ *for parameters* $m, B \geq 2$*,* C*, and color distribution* γ *for the* sum-BC*. We then have*

$$T^v_{\mathrm{OB}(n)}(s^{\mathbf{n}(1)}) \leq_{\mathrm{st}} T^v_{\mathrm{OB}(n)}(s^{\mathrm{k}(1)})$$

for all states $s \in \mathscr{S}_{\text{sum-BC}}$*,* $n \in \mathbb{N}$*, and* $v \in V$*.*

Proof. We will show the stronger $T^v_{\mathrm{OB}}(s^{\mathrm{n}(1,c)}) \leq_{\mathrm{st}} T^v_{\mathrm{OB}}(s^{\mathrm{k}(1)})$ for all $c \notin C_1(s)$ by constructing a coupling Markov chain $Z = (X, Y)$ on a state space which is a subset of $\mathscr{S}_{\text{sum-BC}} \times \mathscr{S}_{\text{sum-BC}}$. The first component of Z behaves exactly as OB started in state $s^{\mathrm{n}(1,c)}$ and the second component as OB started in $s^{\mathrm{k}(1)}$.

A state $(s^{\mathrm{n}}, s^{\mathrm{k}})$ of Z that can be reached from the initial state $(s^{\mathrm{n}(1,c)}, s^{\mathrm{k}(1)})$ will always satisfy the invariant

- either $\chi(s^{\mathrm{n}}) \geq \chi(s^{\mathrm{k}})$, $f_1(s^{\mathrm{n}}) = f_1(s^{\mathrm{k}})$, and $C_1(s^{\mathrm{n}}) = C_1(s^{\mathrm{k}})$ or
- $\chi(s^{\mathrm{n}}) = \chi(s^{\mathrm{k}}) + 1$, $f_1(s^{\mathrm{n}}) = f_1(s^{\mathrm{k}})$, and $C_1(s^{\mathrm{n}}) = C_1(s^{\mathrm{k}}) \cup \{c\}$.

Since in both cases $\chi(s^{\mathrm{n}}) \geq \chi(s^{\mathrm{k}})$, the invariant implies

$$\Pr\left[T^v_X(s^{\mathrm{n}(1,c)}) \leq T^v_Y(s^{\mathrm{k}(1)})\right] = 1,$$

so by Strassen's Theorem (see e. g., [20] or Theorem B.2 in the appendix) the stochastic dominance is established.

It remains to describe Z. The initial state is $(s^{\mathrm{n}(1,c)}, s^{\mathrm{k}(1)})$, which obviously satisfies the invariant. Consider any state $(s^{\mathrm{n}}, s^{\mathrm{k}})$ satisfying the invariant. If s^{n} and s^{k} differ at most in the χ-component, then the transitions of Z are such that the same happens in both components, leading to further states satisfying the invariant.

Suppose s^{n} and s^{k} differ also in the C_1-component. The transitions are then determined by the next color c' drawn according to γ as follows:

$$\begin{cases} \left(s^{\mathrm{n},\mathrm{n}(1,c')}, s^{\mathrm{k},\mathrm{n}(1,c')}\right) & c' \notin C_1(s^{\mathrm{n}}) = C_1(s^{\mathrm{k}}) \cup \{c\}, \\ \left(s^{\mathrm{n},\mathrm{k}(1)}, s^{\mathrm{k},\mathrm{n}(1,c')}\right) & c' = c, \\ \left(s^{\mathrm{n},\mathrm{k}(1)}, s^{\mathrm{k},\mathrm{k}(1)}\right) & c' \in C_1(s^{\mathrm{k}}). \end{cases}$$

Note that all the states satisfy the invariant and that the second kind of transition leads to states which differ at most in the χ-component (the other way of reaching such a state is when bin 1 is empty again). Finally, we can verify that these transitions mirror the behavior of the OB chain in each component:

$$\begin{aligned}
\Pr\left[X_{n+1} = s^{\mathrm{n},\mathrm{n}(1,c')} \middle| X_n = s^{\mathrm{n}}\right] &= 1 - \gamma(C_1(s^{\mathrm{n}})), \\
\Pr\left[X_{n+1} = s^{\mathrm{n},\mathrm{k}(1)} \middle| X_n = s^{\mathrm{n}}\right] &= \gamma(C_1(s^{\mathrm{n}})), \\
\Pr\left[Y_{n+1} = s^{\mathrm{k},\mathrm{n}(1,c')} \middle| Y_n = s^{\mathrm{k}}\right] &= 1 - \gamma(C_1(s^{\mathrm{k}})), \\
\Pr\left[Y_{n+1} = s^{\mathrm{k},\mathrm{k}(1)} \middle| Y_n = s^{\mathrm{k}}\right] &= \gamma(C_1(s^{\mathrm{k}})).
\end{aligned}$$

Theorem 4.9 *Let* OB *and* GF *be the* OneBin *and* GreedyFit sum-BC-*Markov chains for fixed parameters m, B, C and color distribution γ. We have for all states $s \in \mathscr{S}_{\text{sum-BC}}$, in particular the initial empty state, and for all $n \in \mathbb{N}_0$ that*

$$\chi\big(\mathrm{GF}(s)_n\big) \leq_{\mathrm{st}} \chi\big(\mathrm{OB}(s)_n\big).$$

5 Concluding Remarks

We introduced a new approach for the probabilistic analysis of online algorithms which is based on the concept of stochastic dominance. We applied this approach to the analysis of online algorithms for bin coloring problems. This analysis explains simulation results much better than the competitive analysis results existing so far and thus resolves an open problem posed in [17].

In order to prove our results, we started by trying several known methods, in particular ones based on monotonicity and couplings. It turns out that these methods are not powerful enough to prove the stochastic dominance between GreedyFit and OneBin. The detailed results of this attempt are reported in the appendix.

Another result of this work is an algorithm based on Theorem B.3 in the appendix that can be used to construct *certificates* for a stochastic dominance relationship between two Markov chains. Using this algorithm, we could investigate another bin coloring algorithm called FixedColors, which uses a fixed color-to-bin assignment to assign the items to bins based on their colors. Observe that FixedColors is somewhere between OneBin and GreedyFit: it takes advantage of all available bins, but uses a static assignment rule of colors to the bins. Not surprisingly, we observed in simulations that the performance of this algorithm is "in-between" OneBin and GreedyFit. The algorithm was able to find certificates for stochastic dominance for all parameters *m*, *B*, and *C* we tried, assuming a uniform color distribution. The results are reported in Table 1, which shows the smallest number l of "lookahead" steps needed to obtain a certificate, i. e., $l = 2$ means that the construction essentially considers two items at once. For full details, see the appendix. It is interesting to note that for $B = 2$ and for GreedyFit against OneBin $l = 1$ is sufficient, indicating that the "proof" is particularly simple. In fact, the case $B = 2$ can in general be dealt with by a straightforward coupling argument and GreedyFit against OneBin admits the proofs presented in this paper. The fact that the remaining cases require $l > 1$ is an indication that more refined proof techniques are needed.

For the future it is interesting to investigate such refined proof techniques, maybe considering weaker stochastic orders. Similar techniques might also apply for other combinatorial online problems like bin packing or paging, see [11] for some first steps in this direction.

Parameters	GF vs. OB	FC vs. OB	GF vs. FC
$m = 2, B = 2, C = 4$	1	1	1
$m = 2, B = 3, C = 6$	1	4	2
$m = 2, B = 4, C = 8$	1	6	8
$m = 2, B = 5, C = 10$	1	8	9
$m = 2, B = 6, C = 12$	1	10	9
$m = 2, B = 7, C = 14$	1	12	11
$m = 3, B = 2, C = 6$	1	1	1
$m = 3, B = 3, C = 9$	1	4	4
$m = 3, B = 4, C = 12$	1	6	6
$m = 3, B = 5, C = 15$	1	8	6
$m = 4, B = 2, C = 8$	1	1	1
$m = 4, B = 3, C = 12$	1	4	3
$m = 4, B = 4, C = 16$	1	6	4
$m = 5, B = 2, C = 10$	1	1	1
$m = 5, B = 3, C = 15$	1	4	2
$m = 5, B = 4, C = 20$	1	6	3
$m = 3, B = 5, C = 6$	1	1	1
$m = 3, B = 5, C = 9$	1	6	5
$m = 3, B = 5, C = 12$	1	8	6
$m = 3, B = 5, C = 15$	1	8	6
$m = 3, B = 5, C = 18$	1	8	6
$m = 3, B = 5, C = 21$	1	8	6
$m = 3, B = 5, C = 24$	1	8	7
$m = 3, B = 5, C = 27$	1	8	7
$m = 3, B = 5, C = 30$	1	8	8
$m = 3, B = 5, C = 33$	1	8	8
$m = 3, B = 5, C = 36$	1	8	8
$m = 3, B = 5, C = 39$	1	8	8
$m = 3, B = 5, C = 42$	1	8	8
$m = 3, B = 5, C = 45$	1	8	10

Table 1: Results of our certificate-construction algorithm. The reported number denotes the smallest look-ahead for which certificates proving stochastic dominance via Theorem B.3 could be constructed.

References

[1] Angelopoulos, S., Dorrigiv, R., López-Ortiz, A.: On the separation and equivalence of paging strategies. In: SODA 2007, pp. 229–237 (2007)

[2] Angelopoulos, S., Schweitzer, P.: Paging and list update under bijective analysis. In: Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms, pp. 1136–1145 (2009)

[3] Applegate, D.L., Cook, W., Dash, S., Espinoza, D.G.: Exact solutions to linear programming problems. Oper. Res. Lett. **35**(6), 693–699 (2007)

[4] Barney, G.C.: Elevator Traffic Handbook: Theory and Practice. Taylor and Francis (2002)

[5] Becchetti, L., Leonardi, S., Marchetti-Spaccamela, A., Schäfer, G., Vredeveld, T.: Average case and smoothed competitive analysis for the multi-level feedback algorithm. Math. Oper. Res. **31**(1), 85–108 (2006)

[6] Borodin, A., El-Yaniv, R.: Online Computation and Competitive Analysis. Cambridge University Press (1998)

[7] Boyar, J., Favrholdt, L.M.: The relative worst order ratio for online algorithms. ACM Trans. Algorithms **3**(2), Article 22 (2007)

[8] Daley, D.J.: Stochastically monotone Markov chains. Z. Wahrsch. Verw. Gebiete **10**, 305–317 (1968)

[9] Doisy, M.: A coupling technique for stochastic comparison of functions of Markov processes. Journal of Applied Mathematics & Decision Sciences **4**(1), 39–64 (2000). Available at `http://www.hindawi.com/GetRegularIssueArticles.aspx?journal=JAMDS&volume=4`

[10] GNU GMP – the GNU multiple precision arithmetic library. Available at `http://gmplib.org/gmp-man-4.2.1.pdf` (2006)

[11] Hiller, B.: Online optimization: Probabilistic analysis and algorithm engineering. Ph.D. thesis, TU Berlin (2009)

[12] Hiller, B., Vredeveld, T.: Probabilistic analysis of online bin coloring algorithms via stochastic comparison. In: Proceedings of the 16th Annual European Symposium on Algorithms, *Lecture Notes in Computer Science*, vol. 5193, pp. 528–539. Springer (2008)

[13] Hiller, B., Vredeveld, T.: Probabilistic alternatives for competitive analysis. Computer Science – Research and Development **27**(3), 189–196 (2012)

[14] Kalyanasundaram, B., Pruhs, K.: Speed is as powerful as clairvoyance. J. ACM **47**(4), 617–643 (2000)

[15] Karlin, A.R., Manasse, M.S., Rudolph, L., Sleator, D.D.: Competitive snoopy caching. Algorithmica **3**(1-4), 79–119 (1988)

[16] Koutsoupias, E., Papadimitriou, C.: Beyond competitive analysis. In: Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science, pp. 394–400 (1994)

[17] Krumke, S.O., de Paepe, W.E., Stougie, L., Rambau, J.: Online bin coloring. In: F.M. auf der Heide (ed.) Proceedings of the 9th Annual European Symposium on Algorithms, *Lecture Notes in Computer Science*, vol. 2161, pp. 74–84 (2001)

[18] Krumke, S.O., de Paepe, W.E., Stougie, L., Rambau, J.: Bincoloring. Theoret. Comput. Sci. **407**(1-3), 231–241 (2008)

[19] Lin, M., Lin, Z., Xu, J.: Almost optimal solutions for bin coloring problems. In: Proceedings of the 16th International Symposium on Algorithms and Computation, *Lecture Notes in Computer Science*, vol. 3827, pp. 82–91. Springer (2005)

[20] Lindvall, T.: Lectures on the coupling method. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons (1992)

[21] ben Mamoun, M., Bušić, A., Fourneau, J.M., Pekergin, N.: Increasing convex monotone Markov chains: Theory, algorithms, and applications. In: A.N. Langville, W.J. Stewart (eds.) MAM 2006: Markov Anniversary Meeting, pp. 189–210. Boson Books (2006)

[22] Mitzenmacher, M.: Bounds on the greedy routing algorithm for array networks. J. Comput. System Sci. **53**, 317–327 (1996)

[23] Müller, A., Stoyan, D.: Comparison Models for Stochastic Models and Risks. John Wiley & Sons (2002)

[24] de Paepe, W.E.: Complexity results and competitive analysis for vehicle routing problems. Technische Universiteit Eindhoven (2002). Ph. D. Thesis

[25] Scharbrodt, M., Schickinger, T., Steger, A.: A new average case analysis for completion time scheduling. J. ACM pp. 121–146 (2006)

[26] Shaked, M., Shanthikumar, J.G.: Stochastic orders and their applications. Probability and mathematical statistics. Academic Press (1994)

[27] Sleator, D.D., Tarjan, R.E.: Amortized efficiency of list update and paging rules. Comm. ACM **28**(2), 202–208 (1985)

A Further comparison methods for Markov chains

Some standard comparison methods for valued Markov chains are known from the literature, see e.g., [23]. It turned out, however, that these results were not strong enough to explain the behavior of GREEDYFIT and ONEBIN. In this appendix, we discuss several general methods for proving stochastic dominance between Markov chains and show why they are not sufficient for our purposes.

A.1 Monotonicity-based Methods

The first comparison result for Markov chains was provided by Daley [8]. It applies to the case that the state space $\mathscr{S}$ is a subset of $\mathbb{N}_0$ and the valuation function χ is the identity, i. e., the value of a state is the number of the state itself. Daley's result is based on the notion of $\leq_{\text{st}}$-monotonicity. A transition matrix P of a Markov chain on state space $\mathscr{S} \subseteq \mathbb{N}_0$ is called $\leq_{\text{st}}$-*monotone*, if

$$s_1, s_2 \in \mathscr{S}, s_1 \leq s_2 \quad \Longrightarrow \quad P(s_1, \cdot) \leq_{\text{st}} P(s_2, \cdot). \tag{7}$$

For a probability distribution λ and a transition matrix P on $\mathscr{S}$ we denote by λP the probability distribution resulting from starting with λ and doing a one-step transition according to P (this probability is given by the vector-matrix product). Daley proved that (7) is equivalent to requiring that for all distributions λ and μ over $\mathscr{S}$ we have

$$\lambda \leq_{\mathrm{st}} \mu \quad \Longrightarrow \quad \lambda P \leq_{\mathrm{st}} \mu P.$$

Theorem A.1 (see [8]) *If the Markov chains X and Y with transition matrices P and Q and state space $\mathscr{S} \subseteq \mathbb{N}_0$ satisfy*

1. $X_0 \leq_{\mathrm{st}} Y_0$,
2. $P(s,\cdot) \leq_{\mathrm{st}} Q(s,\cdot)$ *for all* $s \in \mathscr{S}$,
3. *at least one of the matrices P and Q is $\leq_{\mathrm{st}}$-monotone,*

then this implies $X_n \leq_{\mathrm{st}} Y_n$ for all $n \in \mathbb{N}_0$.

Obviously, this result is not sufficient to compare valued Markov chains with more complex countable state spaces, since in general different states may have the same value. In this case, χ induces a partial order $\leq_\chi$ on the state space $\mathscr{S}$ by

$$s \leq_\chi s' \quad :\Longleftrightarrow \quad \chi(s) \leq \chi(s'). \tag{8}$$

Using the theory of *integral stochastic orders* (see e. g., [23]) $\leq_{\mathrm{st}}$ can be generalized to partially ordered spaces. We will briefly recapture the notions and results needed to apply monotonicity techniques to valued Markov chains.

Definition A.2 Let $\mathscr{F}$ be a class of functions from $\mathscr{S}$ to $\mathbb{R}$. The class $\mathscr{F}$ induces a stochastic order $\leq_{\mathscr{F}}$ among $\mathscr{S}$-valued random variables X and Y by

$$X \leq_{\mathscr{F}} Y \quad :\Longleftrightarrow \quad \mathbb{E}\left[f(X)\right] \leq \mathbb{E}\left[f(Y)\right] \quad \forall f \in \mathscr{F}. \tag{9}$$

A stochastic order arising in this way is called an *integral stochastic order* with *generator $\mathscr{F}$*.

Definition A.2 defines a very broad class of stochastic orders that includes the stochastic dominance order on arbitrary partially ordered sets. Consider a partial order $\prec$ on $\mathscr{S}$. A function $f\colon \mathscr{S} \to \mathbb{R}$ is called nondecreasing w. r. t. $\prec$ if it satisfies $f(s) \leq f(s')$ whenever $s \prec s'$ for all $s, s' \in \mathscr{S}$. Then the stochastic dominance order $\leq_{\mathrm{st}}$ induced by $\prec$ is generated via (9) by the set of all functions that are nondecreasing w. r. t. $\prec$. The stochastic dominance order $\leq_{\mathrm{st}}$ used so far and defined by inequality (1) is just the special case $\mathscr{S} = \mathbb{R}$ and $\prec = \leq$.

It is easier to check (9) if the class $\mathscr{F}$ is small and has a simple structure. The stochastic dominance order $\leq_{\mathrm{st}}$ induced by a partial order $\prec$ has also a nice small generator. A subset $S \subseteq \mathscr{S}$ is called a *$\prec$-increasing set* if $s \in S$ and $s \prec s'$ imply $s' \in S$. It is not hard to see that $\leq_{\mathrm{st}}$ induced by $\prec$ on $\mathscr{S}$ is generated by the indicator functions of all $\prec$-increasing subsets of $\mathscr{S}$, which are of course $\prec$-nondecreasing functions.

The theorem of Daley can be generalized to any integral stochastic order by generalizing the equivalent characterization of $\leq_{\mathrm{st}}$-monotonicity. We call a stochastic matrix P over $\mathscr{S}$ *$\leq_{\mathscr{F}}$-monotone*, if for all distributions λ and μ over $\mathscr{S}$ we have

$$\lambda \leq_{\mathscr{F}} \mu \quad \Longrightarrow \quad \lambda P \leq_{\mathscr{F}} \mu P.$$

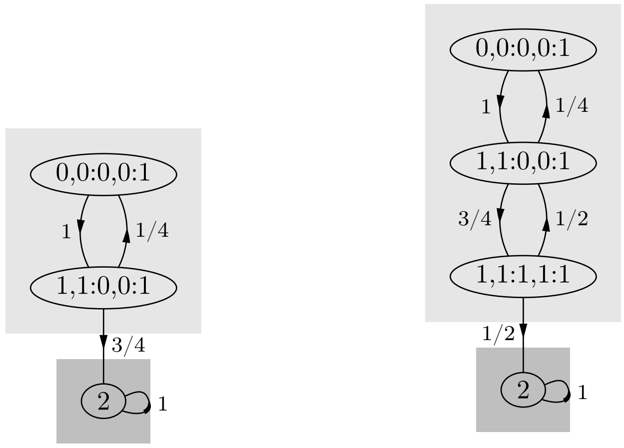

(a) ONEBIN Markov chain OB$(2,2,4)$.

(b) GREEDYFIT Markov chain GF$(2,2,4)$.

Figure 2: Small example Markov chains for ONEBIN and GREEDYFIT for the parameters $m = B = 2$ and $C = 4$. The lightly shaded states are the states of maximum colorfulness 1, the darker one is state $s_{\max}$ with maximum colorfulness 2.

Theorem A.3 (see [23]) *If the Markov chains X and Y with transition matrices P and Q satisfy*

1. $X_0 \leq_{\mathscr{F}} Y_0$,
2. $P(s,\cdot) \leq_{\mathscr{F}} Q(s,\cdot)$ *for all* $s \in \mathscr{S}$,
3. *at least one of the matrices P and Q is $\leq_{\mathscr{F}}$-monotone,*

then this implies $X_n \leq_{\mathscr{F}} Y_n$ for all $n \in \mathbb{N}_0$.

Note that the second requirement $P(s,\cdot) \leq_{\mathscr{F}} Q(s,\cdot)$ can be interpreted as "P is state-wise better than Q". Applied to online algorithms, we can view monotonicity as an additional property that makes an algorithm superior to all algorithms which are not better in any state. This criterion might thus provide guidance for constructing good online algorithms. Unfortunately, it turns out that this result is not strong enough to explain e. g., the stochastic dominance observed between GREEDYFIT and ONEBIN. In the remainder of this section we will show this.

We say that an online bin coloring algorithm is a *partitioning algorithm*, if it uses at most one bin for the same color at any point in time. Observe that ONEBIN, FIXEDCOLORS, and GREEDYFIT are all partitioning algorithms. Assuming a uniform color distribution, the valued Markov chain model for max-BC presented in Section 3 may be simplified for partitioning algorithms as follows: Instead of keeping the *set of colors C_i* in bin i, it is sufficient to keep track of the *number of colors c_i*. For convenience, we assume that $\chi \geq 1$ and we collapse all states with maximum colorfulness to a single state $s_{\max}$ since we are only interested in the colorfulness distribution. Moreover, we break the symmetry by requiring that $\big(f_i(s), c_i(s)\big)$ is lexicographical not smaller than $\big(f_{i+1}(s), c_{i+1}(s)\big)$ in order to reduce the size of the state space. Two small examples are given in Figure 2.

Let $\mathrm{GF} = (\mathrm{GF}_n)_{n\in\mathbb{N}_0}$ and $\mathrm{OB} = (\mathrm{OB}_n)_{n\in\mathbb{N}_0}$ be the valued Markov chains modelling the working of GREEDYFIT and ONEBIN, respectively, for some fixed parameters $m \geq 2$, B, and number of colors C and denote by P_{OB} and P_{GF} the corresponding transition matrices. Moreover, let $\leq_{\chi-\mathrm{st}}$ be the stochastic dominance order $\leq_{\mathrm{st}}$ induced by the partial order $\leq_\chi$ on $\mathscr{S}$. The following result states that GF and OB satisfy requirements 1 and 2 of Theorem A.3 w. r. t. $\leq_{\chi-\mathrm{st}}$.

Proposition A.4 *Let* GF *and* OB *be valued Markov chains as above.*

1. $\mathrm{GF}_0 \leq_{\chi-\mathrm{st}} \mathrm{OB}_0$.
2. $P_{\mathrm{GF}}(s,\cdot) \leq_{\chi-\mathrm{st}} P_{\mathrm{OB}}(s,\cdot)$ *for every state* $s \in \mathscr{S}$.

Proof. The first assertion follows from the fact that both GF and OB start in the same state, i. e., $\mathrm{GF}_0 = \mathrm{OB}_0$.

To establish the second assertion, consider a state $s \in \mathscr{S}$ and let S^{GF} and S^{OB} be random states distributed according to $P_{\mathrm{GF}}(s,\cdot)$ and $P_{\mathrm{OB}}(s,\cdot)$, respectively. We need to show $\mathbb{E}\left[f(S^{\mathrm{GF}})\right] \leq \mathbb{E}\left[f(S^{\mathrm{OB}})\right]$ for all $\leq_\chi$-nondecreasing functions f. It is sufficient to consider the indicator functions $\mathbb{1}_M$ of all $\leq_\chi$-increasing sets M only. As the valuation χ is increasing and can increase at most by one in each step, we have $\mathbb{E}\left[\mathbb{1}_M(S^{\mathrm{GF}})\right] = \mathbb{E}\left[\mathbb{1}_M(S^{\mathrm{OB}})\right]$ whenever there is a state $s' \in M$ with $\chi(s') \leq \chi(s)$ or all states in M have valuation at least $\chi(s)+2$. We therefore can restrict ourselves to the remaining sets M only. Finally, we can assume that state s is reachable in OB, since otherwise we can modify $P_{\mathrm{OB}}(s,\cdot)$ to be the same as $P_{\mathrm{GF}}(s,\cdot)$ without affecting the observed behavior of the OB Markov chain.

A state s reachable by OB is of the form $s = (f_1, c_1, 0, \ldots, 0, \chi)$. Consider a set $M \subseteq \mathscr{S}$ where the lowest valuation of any state is $\chi + 1$. If $f_1 = c_1 = \chi = 0$, we trivially have $\mathbb{E}\left[\mathbb{1}_M(S^{\mathrm{GF}})\right] = \mathbb{E}\left[\mathbb{1}_M(S^{\mathrm{OB}})\right]$ since both algorithms behave identically. In any other case, GF in state s will put an item with a new color in the second bin, thus reaching a state with valuation χ in any case. However, OB will by definition always use the first bin, thus possibly reaching a state in M. Hence, we have

$$\mathbb{E}\left[\mathbb{1}_M(S^{\mathrm{GF}})\right] = \Pr\left[S^{\mathrm{GF}} \in M\right] = 0 \leq \Pr\left[S^{\mathrm{OB}} \in M\right] = \mathbb{E}\left[\mathbb{1}_M(S^{\mathrm{OB}})\right]$$

and the claim follows.

All we still need to do in order to prove $\chi(\mathrm{GF}_n) \leq_{\mathrm{st}} \chi(\mathrm{OB}_n)$ by applying Theorem A.3 is to show that P_{OB} or P_{GF} is $\leq_{\chi-\mathrm{st}}$-monotone. However, this is not necessary and we can get along with a weaker result. Instead, it would be sufficient to show that one of P_{OB} and P_{GF} is monotone w. r. t. some integral order $\leq_{\mathscr{F}}$, where $\mathscr{F}$ is a set of indicator functions of $\leq_\chi$-increasing sets including the *level sets* $\mathscr{S}_{\geq v}$ defined by

$$\mathscr{S}_{\geq v} := \{s \in \mathscr{S} \mid \chi(s) \geq v\}$$

for each valuation v. To see this, suppose $\mathscr{F}$ contains the indicator functions g_v of the level sets $\mathscr{S}_{\geq v}$ and one of the transition matrices is $\leq_{\mathscr{F}}$-monotone. Since requirements 1 and 2 of Theorem A.3 are fulfilled for $\leq_{\chi-\mathrm{st}}$, they are also fulfilled for $\leq_{\mathscr{F}}$, as $\mathscr{F}$ is a set of $\leq_\chi$-nondecreasing functions. We can thus apply Theorem A.3 yielding

$$\begin{aligned}
\Pr[\chi(\mathrm{GF}_n) \geq v] &= \Pr[\mathrm{GF}_n \in \mathscr{S}_{\geq v}] \\
&= \mathbb{E}[g_v(\mathrm{GF}_n)] \\
&\leq \mathbb{E}[g_v(\mathrm{OB}_n)] && \text{by Theorem A.3 and Definition A.2} \\
&= \Pr[\chi(\mathrm{OB}_n) \geq v]
\end{aligned}$$

for all valuations v, which is equivalent to $\chi(\mathrm{GF}_n) \leq_{\mathrm{st}} \chi(\mathrm{OB}_n)$.

However, we can show that this approach does not work. To do this we need some more machinery.

Definition A.5 The *maximal generator* $\mathscr{R}_{\mathscr{F}}$ of an integral stochastic order $\leq_{\mathscr{F}}$ is the set of all functions for which $X \leq_{\mathscr{F}} Y$ implies $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$.

The maximal generator is useful to characterize $\leq_{\mathscr{F}}$-monotone transition matrices. We will use the structure of the maximal generator to show that monotonicity arguments are not suitable for establishing stochastic dominance between online bin coloring algorithms.

Theorem A.6 (see [23]) *A stochastic matrix P is $\leq_{\mathscr{F}}$-monotone if and only if for every $f \in \mathscr{F}$ the function f^P defined by*

$$f^P(s) := \mathbb{E}[f(P(s,\cdot))]$$

is an element of the maximal generator $\mathscr{R}_{\mathscr{F}}$.

Theorem A.7 (see [23]) *If $\mathscr{F} \subseteq \mathscr{G} \subseteq \mathscr{R}_{\mathscr{F}}$, and $\mathscr{G}$ is a convex cone containing the constant functions and being closed under pointwise convergence, then $\mathscr{G} = \mathscr{R}_{\mathscr{F}}$.*

Corollary A.8 *Consider an integral stochastic order $\leq_{\mathscr{F}}$ defined by a* finite *class $\mathscr{F}$ of functions from $\mathscr{S}$ to $\mathbb{R}$. The maximal generator of $\leq_{\mathscr{F}}$ is the convex cone generated by $\mathscr{F}$.*

Proof. It is easy to see that any function $f \in \operatorname{cone}\mathscr{F}$ satisfies $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$ whenever $X \leq_{\mathscr{F}} Y$. Hence, $\mathscr{F} \subseteq \operatorname{cone}\mathscr{F} \subseteq \mathscr{R}_{\mathscr{F}}$. Since $\operatorname{cone}\mathscr{F}$ is finitely generated, it is closed under pointwise convergence and the result follows from Theorem A.7.

The following result shows that the transition matrices of the valued Markov chains for GREEDYFIT and ONEBIN are not $\leq_{\mathscr{F}}$-monotone for stochastic orders of the type described above. Thus monotonicity methods cannot be applied (that easily) to show stochastic dominance results between these algorithms.

Theorem A.9 *Let (X,χ) be a valued Markov chain corresponding to* GREEDYFIT *or* ONEBIN *operating on a color sequence drawn from the uniform distribution on C colors using m bins with capacity $B \geq 3$. Furthermore, let k be the maximum colorfulness attained by the algorithm. If $k \geq 2$, there is no integral stochastic order $\leq_{\mathscr{F}}$ where $\mathscr{F}$ is a (finite) set of indicator functions of $\leq_\chi$-increasing sets including the level sets and the transition matrix P of X is $\leq_{\mathscr{F}}$-monotone.*

Proof. It suffices to consider GREEDYFIT since ONEBIN is obtained by running GREEDYFIT with one bin only.

Assume that P is $\leq_{\mathscr{F}}$-monotone and let g_k be the indicator functions of the level set $\mathscr{S}_{\geq k}$. Denote by $s_1,\ldots,s_l$ all predecessor states of $s_{\max}$. According to the definition of GREEDYFIT, a predecessor state s_j of $s_{\max}$ is of the form $c_i(s_j) = k-1 \geq 1$ for any bin i. Therefore, the probability α that the next color leads to state $s_{\max}$ is the same for all predecessor states and it is strictly positive. Hence we have that $g_k^P(s) = \Pr[\mathscr{S}_{\geq k} \mid s]$ (confer Theorem A.6) is

$$g_k^P(s) = \begin{cases} 0 & s \neq \{s_{\max}, s_1, \ldots, s_l\}, \\ \alpha & s = s_j \quad 1 \leq j \leq l, \\ 1 & s = s_{\max}. \end{cases}$$

Since $g_k \in \mathscr{F}$, P is $\leq_{\mathscr{F}}$-monotone, and cone$\mathscr{F}$ is by Corollary A.8 the maximal generator of $\leq_{\mathscr{F}}$, g_k^P is in cone$\mathscr{F}$ due to Theorem A.6. Thus for any s_j, $1 \leq j \leq l$, there exists a subset $M(s_j) \ni s_j$ of all $s_{\max}$-predecessor states, such that the indicator function $f_{s_j} := \mathbb{1}_{M(s_j) \cup \{s_{\max}\}}$ of $M(s_j) \cup \{s_{\max}\}$ is in $\mathscr{F}$, as such a subset is needed in any conic combination of indicator functions of $\leq_\chi$-increasing sets giving g_k^P.

Now consider the special $s_{\max}$-predecessor state $s = (B-1, k-1, \dots, B-1, k-1, k-1)$ and its predecessor $s' = (B-1, k-1, \dots, B-1, k-1, B-2, k-2, k-1)$, in which one color less than in s is known. We can now compute $f_s^P(s) = \Pr[M(s) \cup \{s_{\max}\} \mid s] = \alpha$, as no state in $M(s)$ can be reached from s. Moreover, $f_s^P(s') = \Pr[M(s) \cup \{s_{\max}\} \mid s'] = \alpha + 1/C$, since s is the only state in $M(S) \cup \{s_{\max}\}$ that can be reached from s'. We thus derived $f_s^P(s') > f_s^P(s)$, but $s' \leq_\chi s$, so f_s^P is not a $\leq_\chi$-nondecreasing function, thus not in cone$\mathscr{F}$. Therefore, P cannot be $\leq_{\mathscr{F}}$-monotone by Theorem A.6.

A.2 Coupling-based Methods

Doisy [9] proposed the following Markov chain comparison criterion, which does not require monotonicity. Recall that given a valued Markov chain (X, χ) and a state $x \in \mathscr{S}$ we use the notation $X(x)$ for the random successor state of state x.

Theorem A.10 (see [9]) *Consider valued Markov chains (X, ϕ) and (Y, ψ) on countable state spaces $\mathscr{S}_X$ and $\mathscr{S}_Y$, respectively. Suppose we have $\phi(X_0) \leq_{\text{st}} \psi(Y_0)$ and*

$$\forall x \in \mathscr{S}_X, y \in \mathscr{S}_Y : \phi(x) \leq \psi(y) \quad \Longrightarrow \quad \phi(X(x)) \leq_{\text{st}} \psi(Y(y)).$$

This implies $\phi(X_n) \leq_{\text{st}} \psi(Y_n)$ for all $n \in \mathbb{N}_0$.

Unfortunately, this criterion is too weak for our purposes, too. To see this, consider the Markov chains of ONEBIN vs. GREEDYFIT in the case $m = B = 2$ and $C = 4$ (cf. Figure 2). For this case, the transition matrices are

$$P_{\text{OB}} = \left(\begin{array}{cc|c} 0 & 1 & 0 \\ 1/4 & 0 & 3/4 \\ \hline 0 & 0 & 1 \end{array}\right), \qquad P_{\text{GF}} = \left(\begin{array}{ccc|c} 0 & 1 & 0 & 0 \\ 1/4 & 0 & 3/4 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ \hline 0 & 0 & 0 & 1 \end{array}\right),$$

where the lines separate the states of colorfulness 1 from the state of colorfulness 2. Let $x = (0,0,1)$ (row 1) be the state of the ONEBIN chain and $y = (1,1,1,1,1)$ (row 3) be the state of the GREEDYFIT chain. Since $\chi(x) = \chi(y) = 1$ we have $\chi(x) \leq \chi(y)$. On the other hand, we have $P_{\text{OB}}(x, \cdot) = (0,1,0)$ and $P_{\text{GF}}(y, \cdot) = (0, 1/2, 0, 1/2)$. Thus $\chi(\text{OB}(x)) = 1$ with probability 1, whereas $\chi(\text{GF}(y)) = 2$ with probability 1/2, so $\chi(P_{\text{GF}}(x, \cdot)) \not\leq_{\text{st}} \chi(P_{\text{OB}}(y, \cdot))$. Hence, the condition of Theorem A.10 is not fulfilled.

B Computer proofs for stochastic dominance relations between bin coloring algorithms

Interestingly, the proof of Doisy's criterion lends itself to a generalization. To describe this, we need the notion of a coupling that is well-known in probability theory (and in some areas of algorithmics as well).

Definition B.1 (Coupling) Let X and Y be $\mathscr{S}$-valued random variables on probability spaces $(\Omega_1, \mathscr{A}_1, P_1)$ and $(\Omega_2, \mathscr{A}_2, P_2)$, respectively. A *coupling of X and Y* is an $\mathscr{S} \times \mathscr{S}$-valued random variable $Z = (\tilde{X}, \tilde{Y})$ on some probability space $(\tilde{\Omega}, \tilde{\mathscr{A}}, \tilde{P})$ with

1. $\tilde{X} \sim X$ i. e., $\tilde{P}\left[\tilde{X} \leq x, \tilde{Y} \text{ arbitrary }\right] = P[X \leq x]$ and

2. $\tilde{Y} \sim Y$.

Couplings are a powerful concept allowing to compare distributions with each other. The following result, which is sometimes refered to as the Theorem of Strassen, gives a well-known connection between stochastic dominance and couplings, see e. g., [20]. It states that stochastic dominance is almost as strong as a pointwise comparison and it is the deeper reason for some of the nice properties of the stochastic dominance order. For instance, the first two statements of Theorem 3.1 can easily be derived from it.

Theorem B.2 (e. g., [20]) *For two random variables X and Y the following are equivalent:*

1. $X \leq_{\mathrm{st}} Y$
2. *There is a coupling $Z = (\tilde{X}, \tilde{Y})$ of X and Y such that* $\Pr\left[\tilde{X} \leq \tilde{Y}\right] = 1$.

Remark For an arbitrary finite set S, consider a uniformly distributed S-valued random variable Σ. Suppose we know $f_1(\Sigma) \leq_{\mathrm{st}} f_2(\Sigma)$ for two real-valued functions f_1 and f_2. The existence of a coupling between $f_1(\Sigma)$ and $f_2(\Sigma)$ according to Theorem B.2 is in this case equivalent to the existence of a bijective mapping $\pi\colon S \to S$ with $f_1(\sigma) \leq f_2(\pi(\sigma))$ for any $\sigma \in S$, which can be obtained by the construction in the proof. Hence, stochastic dominance analysis includes bijective analysis as a special case.

The following theorem gives a sufficient criterion for a pairwise comparison between online bin coloring algorithms. It will serve as the theoretical basis for the computer proofs of stochastic dominance between GREEDYFIT and ONEBIN. For a Markov chain $X = (X_n)_{n \in \mathbb{N}_0}$ the l-step Markov chain $X^l = (X^l_n)_{n \in \mathbb{N}_0}$ is defined by

$$X^l_n := X_{ln} \quad \forall n \in \mathbb{N}_0.$$

If P is the transition matrix of X, the transition matrix of X^l is P^l.

Theorem B.3 *Let (X, ϕ) and (Y, ψ) be valued Markov chains on countable state spaces $\mathscr{S}_X$ and $\mathscr{S}_Y$, respectively. Suppose there is a set $S \subseteq \{(x,y) \in \mathscr{S}_X \times \mathscr{S}_Y \mid \phi(x) \leq \psi(y)\}$ and an $l \geq 1$ s. t.*

1. *For $0 \leq i < l$ there exists a coupling $(\tilde{X}_i, \tilde{Y}_i)$ of X_i and Y_i with*

$$\Pr\left[(\tilde{X}_i, \tilde{Y}_i) \in S\right] = 1$$

and

2. *for every $(x,y) \in S$ there exists a coupling $(\tilde{X}^l, \tilde{Y}^l)$ of $X^l(x)$ and $Y^l(y)$ with*

$$\Pr\left[(\tilde{X}^l, \tilde{Y}^l) \in S\right] = 1.$$

Then we have $\phi(X_n) \leq_{\mathrm{st}} \psi(Y_n)$ for all $n \in \mathbb{N}_0$.

Proof. $0 \leq i < l$ Follows directly using Theorem B.2 and Property 1.

$i \geq l$ We construct a coupling $(\tilde{X}_i, \tilde{Y}_i)$ of X_i and Y_i as follows. First observe that the couplings required by 2. define a transition matrix Z of a Markov chain on state space S where the first component of the state evolves as the l-step Markov chain corresponding to X and the second component as that of Y. Now let $i = ql + r$, $0 \leq r < l$. Let p be the probability distribution of $(\tilde{X}_r, \tilde{Y}_r)$ on S and define $(\tilde{X}_i, \tilde{Y}_i)$ as the random state distributed as pZ^q. By construction, $(\tilde{X}_i, \tilde{Y}_i)$ is in S with probability 1 and has marginal distributions X_i and Y_i, so by Theorem B.2 the claim follows.

Observe that the set of couplings, or equivalently, the pair (S, l), assumed by Theorem B.3 constitute a *certificate* for the stochastic dominance between the valued Markov chains (X, ϕ) and (Y, ψ). We now describe how such a certificate can be found using a computer-aided construction.

We assume that the colors are generated according to the uniform distribution and consider the Markov chain model of Section A.1. In particular, we study the valued Markov chains for the algorithms ONEBIN, FIXEDCOLORS, and GREEDYFIT. Figure 3 gives an example of an explicit construction of the set S for $l = 1$ for a FIXEDCOLORS and a GREEDYFIT Markov chain, indicating the obstacles that may occur.

Of course, one can try to algorithmically find a subset S of the product state space $\mathscr{S}_X \times \mathscr{S}_Y$ and an integer l such that the couplings required by Theorem B.3 exist. Our algorithm (see Figure 4) maintains a set $F \subseteq \mathscr{S}_X \times \mathscr{S}_Y$ of forbidden states, which initially contains those pairs of states $(x, y) \in \mathscr{S}_X \times \mathscr{S}_Y$ with $\phi(x) > \psi(y)$. For a given l, we can assume w. l. o. g. that each Markov chain has l initial states having the i-step distribution as successor distribution, for $0 \leq i < l$. Let $x_0, \ldots, x_{l-1}$ and $y_0, \ldots, y_{l-1}$ be those states for the chains X and Y, respectively. Starting from the states $(x_0, y_0), \ldots, (x_{l-1}, y_{l-1})$, we iteratively try to construct successor-state couplings of all states that can be reached from there such that the probability of reaching a forbidden state is zero. The set G is the set of "good" states that, if the construction is successful, constitute the set S. To keep track of the states for which we still need to find successor-state couplings avoiding F we use a queue Q.

This check is accomplished by procedure COUPLING$(X(x), Y(y), F)$ which tries to compute for random states $X(x)$ and $Y(y)$ and a forbidden set $F \subseteq \mathscr{S}_X \times \mathscr{S}_Y$ a coupling $(\tilde{X}, \tilde{Y})$ that avoids states in F. This procedure can be implemented by solving the following Linear Program. Let $p_{x,y}$ be a real variable for the probability to go to state (x, y).

$$
\begin{aligned}
\min \quad & 0 & \\
\text{s. t.} \quad & \sum_{y \in \mathscr{S}_Y} p_{x,y} = \Pr[X(x) = x] & \forall x \in \mathscr{S}_X, \\
& \sum_{x \in \mathscr{S}_X} p_{x,y} = \Pr[Y(y) = y] & \forall y \in \mathscr{S}_Y, \\
& p_{x,y} = 0 & \forall (x, y) \in F, \\
& 0 \leq p_{x,y} \leq 1 & \forall x \in \mathscr{S}_X, y \in \mathscr{S}_Y.
\end{aligned}
$$

If this Linear Program is feasible, COUPLING$(X(x), Y(y), F)$ indicates success and returns the computed probabilities in matrix Z, otherwise it indicates the failure.

If the coupling construction was successful for state (x, y), we need to ensure that we can successfully couple from all states reached via the successor-state coupling. Therefore we add (x, y) to Q and G if it is not in G yet. In the case there is no F-avoiding successor-state coupling from state (x, y), state (x, y) needs to be forbidden,

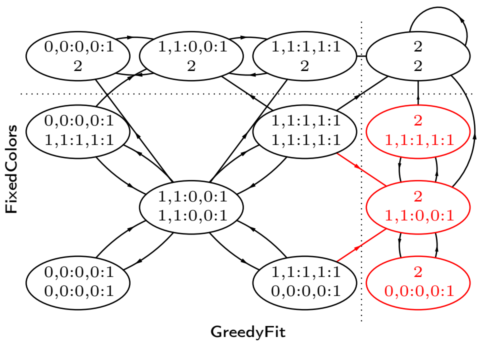

(a) First stage: Product coupling with forbidden states (i. e., states that cannot be part of S) and forbidden transitions into the forbidden states.

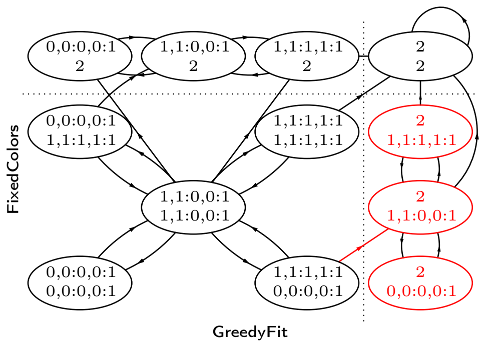

(b) Second stage: The transition from $(1{,}1;1{,}1{:}1 \mid 1{,}1;1{,}1{:}1)$ to $(1{,}1;0{,}0{:}1 \mid 2)$ can be avoided via coupling.

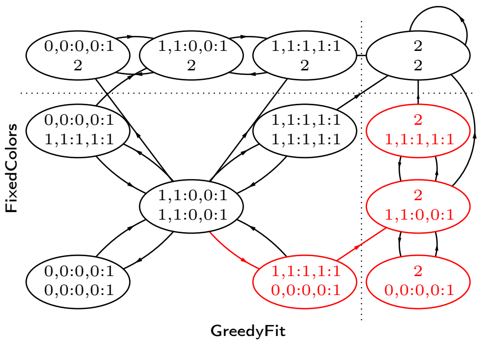

(c) Third stage: The transition from $(0{,}0;0{,}0{:}1 \mid 1{,}1;1{,}1{:}1)$ to $(1{,}1;0{,}0{:}1 \mid 2)$ *cannot* be avoided via coupling, so the state $(0{,}0;0{,}0{:}1 \mid 1{,}1;1{,}1{:}1)$ has to be forbidden as well. The ingoing transition can be avoided via coupling.

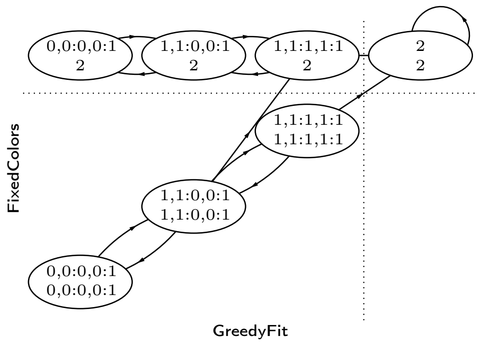

(d) Final stage: Resulting coupling consisting of the set S and the transitions staying in S.

Figure 3: Example construction of a set S that for $l = 1$ shows that GREEDYFIT is superior to FIXEDCOLORS for $m = B = 2$. The upper component of each state is the GREEDYFIT state, the lower the FIXEDCOLORS state.

Input: Valued Markov chains (X,ϕ) and (Y,ψ) and their initial states $x_0,\ldots,x_{l-1}$ and $y_0,\ldots,y_{l-1}$

```
F ← {(x,y) ∈ 𝒮_X × 𝒮_Y | φ(x) > ψ(y)}
Q, G ← ∅
for i = 0,...,l−1 do
    G ← G ∪ {(x_i, y_i)}
    if ¬UPDATEVIACOUPLING(x_i, y_i) then
        stop; construction failed
    end if
end for
while Q ≠ ∅ do
    remove state s = (s_1, s_2) from Q
    if ¬UPDATEVIACOUPLING(s_1, s_2) then
        stop; construction failed
    end if
end while
construction successful, S = G

procedure UPDATEVIACOUPLING(states x and y)
    (success, (X̃, Ỹ)) ← COUPLING(X(x), Y(y), F)
    if success then
        for s′ ∈ {(x,y) ∈ 𝒮_X × 𝒮_Y | Z(x,y) > 0} do
            if s′ ∉ G then
                G ← G ∪ {s′}, Q ← Q ∪ {s′}
            end if
        end for
    else
        F ← F ∪ {(x,y)}, G ← G \ {(x,y)}
        if (x,y) = (x_i, y_i) for some 0 ≤ i < l (initial state) then
            return FALSE                                    ▷ construction failed
        end if
        Q ← Q ∪ {predecessors of (x,y) according to Z}
    end if
    return TRUE
end procedure
```

Figure 4: Algorithm to construct a certificate for stochastic dominance.

too. Once one of the initial states has to be forbidden, it is clear that we cannot avoid entering a forbidden state at all and thus no set S exists.

We implemented the algorithm based on exact arithmetic using the GNU Multiple Precision Arithmetic Library [10] (GMP) and the exact LP solver QSOPT_EX [3] in C++. Table 1 reports our results for the three bin coloring algorithms, in particular the smallest value of l such that our algorithm found appropriate couplings according to Theorem B.3. The correctness of the constructed couplings does not depend on the construction method, but only on the correct implementation of the (simple) routine which checks the conditions of Theorem B.3 and the correctness of the rational arithmetic provided by GMP used in that routine. Assuming this, we get the following result.

Theorem B.4 *Denote by* GF, FC, *and* OB *the valued Markov chains for the* max-BC *for sequences with the uniform color distributions. For the values of m, B, and C given in Table 1, we have*

$$\chi\big(\mathrm{GF}_n(m,B,C)\big) \leq_{\mathrm{st}} \chi\big(\mathrm{FC}_n(m,B,C)\big) \leq_{\mathrm{st}} \chi\big(\mathrm{OB}_n(m,B,C)\big)$$

for all $n \in \mathbb{N}_0$.