

Lower bounds for on-line graph coloring

Magnus M. Halldórsson

15 Asahidai, Tatsunokuchi, Ishikawa, 923-12, Japan

Mario Szegedy

AT&T Bell Labs, 600 Mountain Avenue, Murray Hill, NJ 07974, USA

Abstract

Halldórsson, M.M. and M. Szegedy, Lower bounds for on-line graph coloring, Theoretical Computer Science 130 (1994) 163–174.

An algorithm for vertex-coloring graphs is said to be *on-line* if each vertex is irrevocably assigned a color before later vertices are considered. We show that for every such algorithm there exists a $\log n$ -colorable graph for which the algorithm uses at least $2n/\log n$ colors. This also holds for randomized algorithms, to within a constant factor, against an oblivious adversary.

We then show that various means of relaxing the constraints of the on-line model do not reduce these lower bounds. The features include presenting the input in blocks of up to $\log^2 n$ vertices, recoloring any fraction of the vertices, presorting vertices by degree, and disclosing the adversary's previous coloring.

1. Introduction

1.1. On-line computation

An *on-line algorithm* answers a sequence of requests under the following, informally specified constraints:

- A request must be answered before the next request arrives.
- No information about future requests is available, including the number and the ordering of the requests.
- Each answer is irrevocable.

Once the questioning is over the answers are compared to the answers given by an off-line algorithm, and the algorithm is ranked by the “quality” of its answers relative

Correspondence to: M. Szegedy, AT&T Bell Labs, 600 Mountain Avenue, Murray Hill, NJ 07974, USA.

to an optimal off-line answer sequence. The requests can be assumed to originate from an all-knowing, devious adversary.

There are several important reasons for studying on-line computation. The most commonly cited reason is that it corresponds naturally to the real world: time is unidirectional, past events cannot be reversed, the future is uncertain, and Murphy's Law rules.

Another reason is that on-line algorithms nicely complement many well-studied algorithmic frameworks, such as real-time computation, incremental/dynamic algorithms, prefix solutions of discrete structures, highly recursive computation, and single-pass, greedy, and first-fit algorithms.

A third reason is that on-line computation forms an elegant framework for analyzing algorithms with incomplete information or incomplete access to the input.

1.2. Graph coloring

The problem of *coloring* a graph is that of assigning the vertices to the fewest bins possible so that no two vertices assigned to the same bin are adjacent. In the on-line version of the problem, a vertex is given along with its edges to the previous vertices; the algorithm must irrevocably assign the vertex to a bin before proceeding to the next vertex, but we do not impose any restrictions on the power of the algorithm. The *performance ratio* of a coloring algorithm is the maximum ratio of the number of bins used to the chromatic number (the minimum number of colors required), ranging over all input graphs.

This problem has been much studied, particularly for specific classes of graphs [10, 4, 7]. Lovász et al. [11] gave an algorithm for general graphs that obtains a performance ratio of $O(n/\log^* n)$, slightly improving the trivial bound of n . Vishwanathan [15] gave a randomized algorithm which attains a performance ratio of $O(n/\sqrt{\log n})$ against an oblivious adversary. His algorithm was modified in [5] to improve the performance ratio to $O(n/\log n)$.

In this paper, we prove a $2n/\log^2 n$ lower bound for deterministic on-line graph coloring, which holds to within a constant factor for randomized algorithms. The previous best lower bounds known were $\Omega(\log n)$ for trees [1, 4, 10], and $\Omega(\log^k n)$ for k -colorable graphs, where k is fixed [15].

1.3. Variations of the on-line models

On-line computing places strong constraints on the algorithm; we would like to extend our lower bounds to a more general class of algorithms obtained by weakening some of the restrictions. In the context of the motivating applications, it is in fact natural to relax various conditions.

Take for instance the “answer before next request” condition. The most important property of a real-time algorithm is the ability to respond within a prescribed delay; such a delay may be sufficient to allow for some *look-ahead*: the viewing of a few of the

subsequent requests. Similarly, in a system with a two-level store, the input can be expected to appear in blocks rather than in individual units. And, a single-pass algorithm may have enough memory to keep a limited number of requests in a queue for later processing, which has, for example, been shown to be useful in a server problem with excursions [2].

The “irrevocability” condition can also often be made more flexible. Decisions may be reversible as long as the changes are infrequent and localized. In the case of the bin-packing problem, fewer bins are needed if constant amount of repacking per item is allowed [3]. It is also well known that inputs sorted by nonincreasing item size – a relaxation of the “unknown ordering” principle – require significantly fewer bins [8].

We consider the above variations in the context of the graph-coloring problem, and show them all to yield limited or no improvement. Our lower bound holds to within a constant factor even if the algorithm is given the advantage of $\log^2 n$ -size look-ahead (or blocked input or input buffer), allowed to reassign colors to a constant fraction of the vertices, or allowed to reorder the input based on vertex degrees. This is optimal in the sense that a greater look-ahead/recoloring would trivialize the problem.

All of our results hold in a model that significantly restricts the power of the adversary. The adversary must construct her own coloring on-line and reveal it after the algorithm answers. The results are optimal within that framework. Finally, the constructions also reveal in advance the input length and the number of colors required, both of which the standard framework specifies to be a priori unknown to the algorithm.

The rest of the article is organized as follows. The following section contains definitions of our “transparent” on-line model and related terminology. Section 3 starts with essential properties of the definitions, leading to the main results: lower bounds for deterministic and randomized on-line coloring. Section 4 continues with lower bounds for the various relaxations of the on-line model, and an explanation of the optimality of the basic construction. Some ideas for further work are suggested at the end.

2. Definitions

We consider graphs with an imposed ordering on the vertices $\langle v_1, v_2, \dots, v_n \rangle$. A *pre-neighbor* of a vertex v_i is an adjacent vertex that precedes v_i in the given ordering. The *pre-adjacencies* of v_i , denoted $\text{Adj}^-(v_i)$, is a list of its pre-neighbors. Note that a sequence of all the pre-adjacencies fully specifies the graph. A function f of the vertices is a *proper coloring* if $v_j \in \text{Adj}^-(v_i)$ implies $f(v_j) \neq f(v_i)$, for each v_j, v_i .

Transparent on-line coloring is a combinatorial game between two players: the *adversary* A , and the *algorithm* B . The game consists of an undetermined, but finite, number of request–answer–reply rounds $1, 2, \dots, n$, where round t is as in Fig. 1.

A poses a pair $(v_t, \text{Adj}^-(v_t))$ where $\text{Adj}^-(v_t) \subseteq \{v_1, \dots, v_{t-1}\}$.
 B answers with an integer $\text{Bin}(v_t)$, a proper coloring of v_t .
 A responds with an integer $\text{Col}(v_t)$, a proper coloring of v_t .

Fig. 1. Transparent on-line coloring game.

Without the last step in each round, the game is equivalent to the usual on-line coloring problem. We say that an adversary is *transparent* if, as above, it reveals its decisions following the algorithm. In this paper, we fix in advance the number of rounds n , with the performance evaluation occurring only at the end. We generally describe only the t th round, intending the description to be applicable to each round.

We say that the algorithm colors with *bins* and the adversary with *colors* in order to easily distinguish the two. Both of these are objects from some finite sets, which we associate with the positive integers. The term bin will refer both to the cardinal assigned to vertices, as well as to the set of vertices that have been assigned that bin number. For a vertex v_t , the set $\text{Avail}(v_t)$ of *admissible colors* consists of the colors *not* used by its pre-neighbors.

Define the *hue* of a bin to be the set of colors of the vertices in the bin, $\text{Hue}(b) = \{\text{Col}(v_i) : \text{Bin}(v_i) = b\}$. A *hue collection* H is a set of all the nonempty hues. Both of these are dependent on an implicit point in time.

The term k will denote the number of colors sufficient to properly color the graph constructed. We shall assume in this paper that k is divisible by 16 in order to simplify the presentation. Let $[k]$ denote the set $\{1, 2, \dots, k\}$, and $[k]^{(k/2)}$ the collection of all subsets of $[k]$ of size $k/2$. Finally, let Random be a function that gets as input a finite set, and selects an item with uniform probability.

3. Main course

3.1. Preliminaries

Our task as the adversary is to decide on the adjacencies and the color of each vertex. The adjacencies are determined by first selecting a set of admissible colors, and then applying the following rule: A vertex v_t shall be adjacent to a previous vertex v_i iff the color of v_i does not belong to the set of admissible colors. The idea is that if we make the vertex adjacent to any previous node of a given color it does not hurt to make it adjacent to all nodes of that color. From this definition, we get the following property.

Observation 3.1. $\text{Hue}(\text{Bin}(v_t)) \subseteq \text{Avail}(v_t)$.

Our task is to select a set of admissible colors, and to select one of those as the vertex color, so that the algorithm is bound to introduce many bins. Our measure of

progress is the number of distinct pairs of bin and color values assigned to the vertices, or $|\{(\text{Bin}(v_i), \text{Col}(v_i)): 1 \leq i \leq t\}|$.

To see that this number is actually indicative of the progress made, notice that the number of bins is at least the number of bin/color pairs divided by the maximum number of pairs with the same bin number. But the latter is bounded by the number of admissible colors for the last vertex assigned to that bin, by Observation 3.1, giving us the following bound.

Observation 3.2. If m distinct bin/color pairs have been assigned after round t , then at least $m/(\max_{i \leq t} |\text{Avail}(v_i)|)$ bins have been used.

It is useful to note that if the bin hue is a strict subset of the set of admissible colors, then we are in a situation to make progress, since then the vertex can be assigned a color different from those of other vertices in the bin.

Observation 3.3. If $\text{Col}(v_t) \in \text{Avail}(v_t) - \text{Hue}(\text{Bin}(v_t))$ then progress is made.

3.2. Adversary against deterministic algorithms

Let k be a positive even number, and

$$n = \frac{k}{2} \binom{k}{k/2}$$

be the number of rounds. Let *any* denote a nondeterministic selection operator that gets as input a set and returns some member. Figure 2 gives the adversary's adaptive strategy; set H is the current set of hues.

Theorem 3.4. The performance ratio of any deterministic on-line coloring algorithm is at least $2n/\log^2 n$.

Proof. Each round of the game contributes at most one element to the combined membership of the hues in H . As a result, as long as the game is played, i.e., as long as

$$t \leq \frac{k}{2} \binom{k}{k/2},$$

at least one of the $\binom{k}{k/2}$ sets of size $k/2$ is not a member of H , and can therefore be assigned to $\text{Avail}(v_t)$. Since $\text{Avail}(v_t)$ cannot equal any bin hue (by definition) but must

$$\begin{aligned} \text{Avail}(v_t) &= \text{any}([k]^{(k/2)} - H) \\ \text{Adj}^-(v_t) &= \{v_i: \text{Col}(v_i) \notin \text{Avail}(v_t) \text{ and } i < t\} \\ \text{Col}(v_t) &= \text{any}(\text{Avail}(v_t) - \text{Hue}(\text{Bin}(v_t))) \end{aligned}$$

Fig. 2. Adversary strategy: adaptive construction.

include the hue of the selected bin (by Observation 3.1), the inclusion must be strict, and $\text{Col}(v_t)$ is thus well defined. Each round therefore makes progress (by Observation 3.3), and the number of bins must be at least $n/(k/2)$ (by Observation 3.2). Since the number of colors is at most k , the performance ratio is at least $2n/k^2$, where $k = \log n(1 - o(1))$. $\square$

To appreciate the simplicity of the construction, the reader is encouraged to work through an example.

3.3. Adversary against randomized algorithms

Randomized on-line algorithms can be evaluated in at least three different ways, depending on the power of the adversary in question. The weakest of these, the *oblivious* adversary, must construct the whole input in advance before feeding it to the algorithm. It is against this adversary that the algorithms of [5, 15] are successful.

We show in this section that for any randomized on-line coloring algorithm there exists a k -colorable graph on which the algorithm will use at least expected n/k bins, where $k = O(\log n)$. By Yao's lemma [16], it suffices to show that there exists a distribution of k -colorable graphs for which the average number of colors used by any *deterministic* algorithm is at least n/k .

We construct a distribution of k -colorable graphs, in $n = 2^{k/4}$ rounds as in Fig. 3. Observe that the construction is oblivious, since decisions made by the algorithm, such as bin assignments, are never consulted.

Consider a given choice of $\text{Avail}(v_t)$. The probability that the random color assignment yields a successful round equals

$$|\text{Avail}(v_t) - \text{Hue}(\text{Bin}(v_t))| / |\text{Avail}(v_t)|.$$

Since the graphs are constructed in advance, we cannot assume anything about the actual values of the hues. Instead, we shall show that for a randomly chosen $\text{Avail}(v_t)$ and any fixed collection H of at most n hues, the difference $|\text{Avail}(v_t) - h|$ minimized over all hues h in H can be expected to be large. For a $k/2$ -set p , and a collection H of subsets of $[k]$, define

$$\text{dist}(p, H) = \min_{h \in H} |p - h|.$$

$$\begin{aligned} \text{Avail}(v_t) &= \text{Random}([k]^{(k/2)}) \\ \text{Adj}^-(v_t) &= \{v_i : \text{Col}(v_i) \notin \text{Avail}(v_t) \text{ and } i < t\} \\ \text{Col}(v_t) &= \text{Random}(\text{Avail}(v_t)) \end{aligned}$$

Fig. 3. Adversary strategy: *oblivious construction*.

We have that for any hue collection H ,

$$|\text{Avail}(v_t) - \text{Hue}(\text{Bin}(v_t))| \geq \text{dist}(\text{Avail}(v_t), H).$$

If we now let $\text{Avail}(v_t)$ vary, the probability of a successful round is at least

$$E[\text{dist}(\text{Avail}(v_t), H)] / |\text{Avail}(v_t)|.$$

The performance ratio lower bound follows easily from the following bound on this distance function.

Lemma 3.5. *Let H be any collection of subsets of $[k]$, and let p be a randomly chosen subset of $[k]$ of size $k/2$. If $|H| \leq 2^{k/4}$ and k is large enough, then*

$$E[\text{dist}(p, H)] \geq k/4. \quad (1)$$

Proof. We claim that

$$\Pr[\text{dist}(p, H) \leq s] \leq \binom{k/2+s}{s} |H| / \binom{k}{k/2}. \quad (2)$$

This implies, using Stirling's approximation, that

$$\Pr[\text{dist}(p, H) \leq 0.28k] \leq 1.01^{-k}$$

for $|H| \leq 2^{k/4}$. Now,

$$E[\text{dist}(p, H)] \geq 0.28k \Pr[\text{dist}(p, H) \geq 0.28k] \geq 0.25k$$

for k large enough.

To see that inequality (2) holds, let H' contain the sets in H whose size is at least $k/2 - s$. Each set in H' is a subset of at most $\binom{k/2+s}{s}$ sets of size $k/2$. Therefore, we get

$$\Pr[\text{dist}(p, H') \leq s] \leq \binom{k/2+s}{s} |H'| / \binom{k}{k/2}$$

by summing the probabilities of the sets in H' . But none of the elements of $H - H'$ affect these probabilities, since they are too far away from p . Hence the claim and the lemma. $\square$

Theorem 3.6. *The performance ratio of any randomized on-line coloring algorithm is at least $n/(16 \log^2 n)$.*

Proof. Consider the execution of a fixed deterministic algorithm on the graphs drawn from the distribution constructed above. The probability that a node makes progress is at least

$$\frac{E[\text{dist}(\text{Avail}(v_t), H)]}{|\text{Avail}(v_t)|} \geq \frac{k/4}{k/2} = 1/2.$$

Hence, the expected number of bin/color pairs is at least $n/2$, by linearity of expectation, and the expected number of bins at least n/k by Observation 3.2. Thus the performance guarantee is at least n/k^2 , where $k = 4 \log n$. $\square$

4. Bells and whistles

4.1. Blocked input

One of the caveats of computing on-line is that an answer is often proved wrong or ineffective almost immediately after it is uttered. One hope would be that if just a little bit was known about what's coming up on the horizon, far better decisions could be made. No such luck. We can show that even with moderate visibility, the performance would not improve.

We construct the input in an oblivious manner, in blocks of size $k^2/128$ organized as a sequence of $k/16$ disjoint cliques of size $k/8$. As before, the vertices are adjacent to vertices in previous blocks according to a randomly chosen set of admissible colors.

The formal definition is given in Fig. 4. If vertex v_t is in the b th block and inside this block is in the c th clique, then $t = bk^2/128 + ck/8 + t'$, where $0 \leq t' < k/8$ (blocks and cliques are counted from 0). The block number and clique number of v_t are $\text{blk}_t = b$ and $\text{cliq}_t = c$. Let $\text{cc}(v_t) = \{\text{Col}(v_i) : i < t \text{ and } \text{blk}_i = \text{blk}_t \text{ and } \text{cliq}_i = \text{cliq}_t\}$ represent the colors of the previous vertices in the same clique. Let $n = 2^{k/4}$ be the number of rounds.

Lemma 4.1. *Against any randomized algorithm with blocked input of size up to $k^2/128$, an expected constant fraction of the rounds in fig. 4 will make progress.*

Proof. As before, the probability of success is a function of the difference between the set of admissible colors and the largest valid bin hue. The expected size of the bin hue at the beginning of the look-ahead block is at most $k/4$, by Lemma 3.5, and the number of vertices added since then cannot exceed the independence number of the subgraph in the block, or $k/16$. The number of admissible colors is $|\text{Avail}(v_t) - \text{cc}(v_t)|$ and ranges from $k/2 - (k/8 - 1)$ to $k/2$. The probability of success of a random color assignment is thus at least

$$\frac{|\text{Avail}(v_t) - \text{cc}(v_t)| - k/4 - k/16}{|\text{Avail}(v_t) - \text{cc}(v_t)|} \geq \frac{k/2 - k/4 - k/8 - k/16}{k/2 - k/8} = 1/6. \quad \square$$

$$\text{Avail}(v_t) = \text{Random}([k]^{(k/2)})$$

$$\text{Adj}^-(v_t) = \{v_i : (\text{blk}_i < \text{blk}_t \text{ and } \text{Col}(v_i) \notin \text{Avail}(v_t)) \text{ or } (\text{blk}_i = \text{blk}_t \text{ and } \text{cliq}_i = \text{cliq}_t)\}$$

$$\text{Col}(v_t) = \text{Random}(\text{Avail}(v_t) - \text{cc}(v_t))$$

Fig. 4. Adversary strategy: blocked input.

Given the bound on the success probability of the preceding lemma, the following theorem follows easily. Notice that in our construction k is of order $\log n$, hence the bound holds even if we measure the block size in terms of n .

Theorem 4.2. *The performance ratio of any randomized algorithm, when the input is presented in blocks of size $O(\log^2 n)$, is $\Omega(n/\log^2 n)$.*

The lower bound is the best possible, since there is a simple deterministic method to take advantage of blocks of larger size. Off-line coloring each block of size l optimally (possibly using exponential time) requires at most $\chi(G)\lceil n/l \rceil$ colors, where $\chi(G)$ is the minimum number of colors required to color G , which implies an n/l performance ratio. Our lower bound matches this for any $l = \Omega(\log^2 n)$, to within a constant factor, by padding each input block (adding $l - k^2/128$ dummy vertices). Therefore, in this problem, the effect of blocked input on the performance guarantee is a threshold function.

Similar threshold-like behavior has been shown for the on-line problems of bipartite matching [9] and coloring inductive graphs [7].

4.2. Look-ahead and buffering

We can treat two interesting variations of blocked input similarly. An algorithm is *on-line with look-ahead l* if it bases its answer to request t only on the first $t + l$ requests. And it is *on-line with a buffer of size l* if at step $t + l$ at least t requests have been answered.

It is easy to see that look-ahead is more powerful (as an algorithmic feature) than blocked input, and that an input buffer is more powerful than look-ahead. Also, we can always simulate look-ahead with blocked input of double the block size: follow every block of actual input, with an equivalent amount of dummy input.

We are also able to deal with buffered coloring effectively. Construct blocks twice the size of the buffer; at least one bufferful – half a blockful – must be answered before the end of the block. By inequality (2), the probability that *every* vertex in the block succeeds is high, or at least

$$1 - \binom{k/2 + k/16}{k/16} |H| \bigg/ \binom{k}{k/2} \leq 1 - 1.38^{-k},$$

which is asymptotically 1. Thus, no matter which half of the block the algorithm answers, the expected amount of success is not decreased, and the bound for look-ahead holds also for buffering to within a factor of two.

This result differs from results on other problems, such as the problem of balancing vectors in a metric space, where look-ahead is of no help while buffering yields dramatic improvements [14, p. 69].

4.3. Recolorings

One suggested alternative to look-ahead is to allow the algorithm to “recolor” a portion of the nodes, i.e., to independently reassign them bins at some point in the coloring process. This would be useful if only a few nodes were the cause of a poor coloring while the assignments were overall reasonable.

We find that significant amounts of recolorings are of limited use. The adaptive adversary in Theorem 3.4 makes progress in every round, and thus forces the algorithm to use at least $r/\log^2 r$ colors, when counting only the r vertices never recolored. Thus, unless almost all, $n - o(n)$, of the vertices are recolored, the $\Omega(n/\log^2 n)$ lower bound on the performance ratio still holds. We note, however, that this claim does not apply to our lower bound argument for randomized algorithms as stated.

4.4. Pre-processing: sorting vertices by degree

One hope for a more effective on-line algorithm would be to pre-massage the input into a pliable ordering. The most natural ordering criteria for graphs are those based on the degrees of the vertices. Such strategies have been extensively evaluated both experimentally and analytically in association with the ubiquitous, inherently on-line First-Fit coloring algorithm (e.g. [13]).

We can easily circumvent any such attempt by padding the input so that all vertices will be of the same degree. With $n/(k-1)$ extra vertices for each of the k colors, each original vertex can then be made adjacent to up to n new vertices without destroying the k -colorability property. The randomized constructions can do with even less padding, due to the highly convergent nature of random selection.

4.5. Transparency and optimality

In the construction of Section 3.2, the adversary makes her coloring assignments on-line, and can without harm reveal those decisions following the algorithm’s answers. When given the advantage of this extra information, there is a simple, effective algorithmic strategy: Allocate bins for every nonempty subset of $[k]$. Assign the current vertex to the bin whose hue is a maximal subset of the admissible colors for the vertex. This ensures that no more than $2^k - 1$ bins are used, and in fact $k2^{k-1}$ vertices are required for all of them to be used.

This can be matched precisely, obtaining a minimax value for the game, if, in the deterministic construction, the adversary selects any *minimal* hue (no longer restricting herself to $k/2$ -sets) unoccupied by a bin. The details can be found in [6].

The strong lower bound obtained here for the transparent model is in stark contrast with its ineffectiveness for other problems. For the three-coloring problem, exactly 7 bins are needed, a far cry from the (nevertheless weak) $\Omega(\log^2 n)$ lower bound for the standard model [15]. Similarly, for the k -server problem [12], we can give a simple

algorithm with a performance ratio of 3, for any $k \geq 3$, compared to the lower bound and conjectured upper bound of k for the standard model.

5. Conclusions

We have presented strong lower bounds for on-line graph coloring. Chances are that further lower bounds can be found to bridge the gap that remains in the deterministic case, but not with the methods of this paper.

We have also given matching bounds for several variants of the standard on-line model. Fundamental questions about properties and the applicability of these and other variants are yet to be studied in a general framework. For instance, is the value of look-ahead always a threshold function?

Acknowledgments

The first-mentioned author would like to thank Jaikumar Radhakrishnan and Gábor Tardos for many helpful discussions and insights. An earlier version of this paper, with different results, appeared in [6].

References

- [1] D. Bean, Effective coloration, *J. Symbolic Logic* **41** (1976) 469–480.
- [2] F.R.K. Chung, R.L. Graham and M.E. Saks, A dynamic location problem for graphs, *Combinatorica* **9**(2) (1989) 111–131.
- [3] G. Gambosi, A. Postiglione and M. Talamo, New algorithms for on-line bin packing, in: *Proc. 1st Italian Conf. on Algorithms* (1990) 44–59.
- [4] A. Gyárfás and J. Lehel, On-line and first fit colorings of graphs, *J. Graph Theory* **12** (2) (1988) 217–227.
- ✂ [5] M.M. Halldórsson, Parallel and on-line graph coloring, in: *Proc. 3rd Internat. Symp. on Algorithms and Computation* (Nagoya, Japan, Dec. 1992), Lecture Notes in Computer Science (Springer, Berlin) 61–70.
- ✂ [6] M.M. Halldórsson and M. Szegedy, Lower bounds for on-line graph coloring, in: L.A. McGeoch and D.D. Sleator, eds., *Proc. DIMACS Workshop on On-Line Algorithms* (AMS–ACM, Feb. 1991), Series in Discrete Mathematics and Theoretical Computer Science, Vol. 7.
- [7] S. Irani, On-line coloring inductive graphs, in: *Proc. 31st Ann. IEEE Symp. on Foundations of Computer Science* (Oct. 1990) 470–479.
- [8] D.S. Johnson, A. Demers, J.D. Ullman, M.R. Garey and R.L. Graham, Worst-case performance bounds for simple one-dimensional packing algorithms, *SIAM J. Comput.* **3** (4) (1974) 299–325.
- [9] M.-Y. Kao and S.R. Tate, Online matching with blocked input, *Inform. Process. Lett.* **38** (1991) 113–116.
- [10] H.A. Kierstead and W.T. Trotter, An extremal problem in recursive combinatorics, in: *Proc. 12th Southeastern Conf. on Combinatorics, Graph Theory, and Computing. Congressus Numerantium XXXIII* (1981) 143–153.
- [11] L. Lovász, M. Saks and W.T. Trotter, An online graph coloring algorithm with sublinear performance ratio, *Discrete Math.* **75** (1989) 319–325.

- [12] M.S. Manasse, L.A. McGeoch and D.D. Sleator, Competitive algorithms for on-line problems, *J. Algebra* **11** (1990) 208–230.
- [13] D.W. Matula and L.L. Beck, Smallest-last ordering and clustering and graph coloring algorithms, *J. ACM*, **30** (1983) 417–427.
- [14] J. Spencer, *Ten Lectures on the Probabilistic Method*, Regional Conference Series in Applied Mathematics, Vol. 52 (SIAM, Philadelphia, PA, 1987).
- [15] S. Vishwanathan, Randomized online graph coloring, *J. Algorithms* **13** (1992) 657–669.
- [16] A.C.-C. Yao, Probabilistic computations: Towards a unified measure of complexity, in: *Proc. 18th Ann. IEEE Symp. on Foundations of Computer Science* (1977) 535–542.