DIMACS Series in Discrete Mathematics
and Theoretical Computer Science
Volume 7, 1992

ON-LINE GRAPH COLORING

H. A. KIERSTEAD† and W. T. TROTTER‡

Abstract. We survey recent results concerning on-line graph coloring and announce the following Theorem: For every radius two tree T, the class Forb(T) of graphs which do not induce T is on-line χ-bounded. In particular, the class of co-comparability graphs is on-line χ-bounded.

Key words. on-line, algorithm, graph, coloring

AMS(MOS) subject classifications.

1. Introduction. In this article we report on recent results concerning on-line graph coloring and suggest a few interesting problems for future research. An *on-line graph* is a structure $G^< = (V, E, <)$, where $G = (V, E)$ is a graph and $<$ is a linear ordering of V. We say that $G^<$ is an *on-line presentation* of G. We shall always assume that $V = \{v_1, \ldots, v_n\}$, where $v_i < v_j$ if and only if $i < j$. In particular, $G^<$ has $n = n(G)$ vertices. Then we let $V_i = \{v_j : j \leq i\}$ and $G_i^< = G^<[V_i]$, the on-line subgraph of $G^<$ induced by V_i. If two vertices v and w are adjacent in G, we write $v \sim w$. The *neighborhood* of a vertex v in G is $N(v) = N_G(v) = \{w \in V : v \sim w\}$. An algorithm for coloring the vertices of an on-line graph $G^<$ is said to be *on-line* if the color of a vertex v_i is determined solely by $G_i^<$. Intuitively, the algorithm colors the vertices of $G^<$ one at a time in the externally determined order $v_1, \ldots, v_n$, and at the time a color is irrevocably assigned to the vertex v_i, the algorithm can only see $G_i^<$. A simple, but important example of an on-line algorithm is the algorithm First-Fit, which colors the vertices of G with an initial sequence of the colors $\{1, 2, \ldots, \}$ by assigning to the vertex v_i the least possible color which is not assigned to any vertex of V_{i-1} adjacent to v_i.

The *clique number* and *chromatic number* of G are denoted by $\omega(G)$ and $\chi(G)$, respectively. For an on-line algorithm $\mathbb{A}$ and an on-line graph $G^<$, let $\chi_{\mathbb{A}}(G^<)$ denote the number of colors $\mathbb{A}$ uses to color $G^<$. The *performance function* $\phi_{\mathbb{A}}(k, n/\Gamma)$ of $\mathbb{A}$ over a class of graphs Γ is defined for integers k and n to be the maximum of $\chi_{\mathbb{A}}(G^<)$ over all on-line presentations of k-colorable graphs $G \in \Gamma$ on n vertices. Note that $\phi(k, n/\Gamma)$ is an increasing function; we denote the limit as n goes to infinity of $\phi_{\mathbb{A}}(k, n/\Gamma)$ by $\phi_{\mathbb{A}}(k/\Gamma)$. When Γ is the class of all graphs, we may simply write

1991 *Mathematics Subject Classification.* Primary 05C15, 05C85.

†Department of Mathematics, Arizona State University, Tempe, AZ 85287–1804, USA. Partially supported by Office of Naval Research grant N00014-90-J-1206.

‡Department of Mathematics, Arizona State University, Tempe, AZ 85287–1804, USA. After August 15, 1991: Bell Communications Research, 445 South Street, Morristown, NJ 07960–1910. Partially supported by National Science Foundation grant DMS 89-0241.

This paper is in final form.

$\phi_{\mathbb{A}}(k,n)$. The following elementary theorem, originally phrased in terms of recursive functions, shows that the performance of an on-line coloring algorithm for an arbitrary graph G cannot be bounded above solely in terms of $\chi(G)$.

THEOREM 1.1. *Bean[1]. For every on-line algorithm $\mathbb{A}$ and integer t, there exists an on-line tree $T^<$ such that $\chi_{\mathbb{A}}(T^<) > t$. Moreover, $T^<$ has only 2^t vertices.*

The tree is constructed using a simplification of the Zykov [32] construction of triangle-free t-critical graphs. Arguing inductively, one constructs disjoint on-line trees $T_i^<$, for $i = 1, \dots, t-1$, so that $T_i^<$ has 2^i vertices and the algorithm is forced to use i colors on $T_i^<$. Then it is possible to choose a vertex x_i in $T_i^<$ so that each x_i, $i = 1, \dots, t-1$, has a distinct color. The final vertex v_{2^t} is played adjacent to each x_i and must receive the t-th color.

The situation is even worse for First-Fit. Let B_t be the graph formed from the complete bipartite graph $K_{t,t}$ by removing a perfect matching M_t. Then B_t has $2t$ vertices, but the on-line presentation $B_t^<$ of B_t, where the pairs of vertices matched in M_t are ordered consecutively, forces First-Fit to use t colors.

In §2 we consider the performance function of on-line coloring algorithms for general graphs. In §3 we consider special classes of graphs for which there exist on-line algorithms whose performance can be bounded solely in terms of clique size. In §4 we consider even more special classes of graphs for which the performance of First-Fit can be bounded solely in terms of clique size.

2. Performance bounds for general graphs.

Our first theorem, when combined with Theorem 1.1, shows that $\phi(2,n) = \theta(\log n)$.

THEOREM 2.1. *Lovász, Saks and Trotter[25]. There exists an on-line algorithm $\mathbb{A}$ such that for every on-line 2-colorable graph $G^<$ on n vertices, $\chi_{\mathbb{A}}(G^<) \le 2lg\ n$.*

When a new point v_i is considered there is a unique partition (I_1, I_2) of the component of v_i in $G_i^<$ into independent sets with $v_i \in I_1$. The algorithm $\mathbb{A}$ assigns v_i the least color not assigned to any vertex of I_2. Observe that if $\mathbb{A}$ assigns v_i color $k+2$, then $\mathbb{A}$ must have already assigned $k+1$ to some vertex of I_2 and k to some vertex $v_p \in I_2$. Thus, $\mathbb{A}$ must have assigned k to some vertex $x_q \in I_1$. Since $\mathbb{A}$ assigned v_p and v_q the same color, v_p and v_q are in separate components of $G_r^<$, where $r = \max\{p, q\}$. Thus, by induction, each of these components must have size $2^{k/2}$, and thus $i > 2(2^{k/2}) = 2^{(k+2)/2}$.

Vishwanathan improved the technique used to prove Theorem 1.1 in order to show that $\phi(k,n) = \Omega(\lg^{k-1} n)$, for fixed k.

THEOREM 2.2. *Vishwanathan[30]. For every on-line algorithm $\mathbb{A}$ and all integers k and n, there exists an on-line graph $G^<$ on n vertices such that $\chi(G^<) \le k$ and $(lg\, n/(4k))^{k-1} \le \chi_{\mathbb{A}}(G^<)$.*

For a fixed algorithm, the on-line graph $G^<$ is constructed using a primary induction on k and a secondary induction on n. The key idea is to maintain the strong induction hypothesis that $\mathbb{A}$ can be forced to use $(\lg n/(4k))^{k-1}$ colors on one part of some partition of $G^<$ into k independent sets. Then the primary induction hypothesis can be used to attach a $k-1$ colorable graph to this part so that many new colors must be used. Some additional care must be taken to maintain the strong induction hypothesis.

Using an elegant construction, Szegedy proved:

THEOREM 2.3. *Szegedy[28]. For every on-line algorithm* $\mathbb{A}$ *and integer* k, *there exists an on-line graph* $G^<$ *on* n *vertices such that* $\chi(G^<) \leq k$, $n \leq k2^k$, *and* $\chi_{\mathbb{A}}(G^<) \geq 2^k - 1$.

The only positive result for all k-colorable graphs, with $k \geq 3$, is due to Lovász, Saks, and Trotter. They prove that $\phi(k,n)$ is sublinear in n. Here $\lg^{(k)}$ denotes lg iterated k times.

THEOREM 2.4. *Lovász, Saks and Trotter[25]. There exists an on-line algorithm* $\mathbb{A}$ *such that for every* k*-colorable on-line graph* $G^<$ *on* n *vertices,* $\chi_{\mathbb{A}}(G^<) = O(n\, lg^{(2k-3)}n/lg^{(2k-4)}n)$. *Moreover, there exists an on-line algorithm* $\mathbb{B}$ *such that for every integer* k *and every on-line* k*-colorable graph* $G^<$ *on* n *vertices,* $\chi_{\mathbb{B}}(G^<) \leq (kn/\log^* n)(1+o(1))$.

We give an overview of the algorithm when $k=3$. The algorithm follows the First-Fit rule as long as this results in an assignment of a color at most $t = 100\log\log n$. Let X denote the remaining vertices. Then each $v \in X$ has in its neighborhood a subset consisting of one vertex from each of the first t color classes. Define a sequence $\alpha_1 = 1, \alpha_2, \alpha_3, \ldots$ by $\alpha_{i+1} = (\alpha_i)^2/2$. The set X is partitioned on-line into subsets $X_1, X_2, X_3, \ldots$ so that if $|X_j| \geq s$, then the vertices in X_i have at least $\alpha_s t$ common neighbors. Subject to this restriction, the new vertex is added to the largest possible subset. If it cannot be added to any of the existing nonempty sets in the partition, then it is added to the partition as a singleton. Each set in the partition is 2-colorable and has its own color set which consists of at most $2\lg|X_i|$ colors. Some details remain. First, the number of sets of size s has to be bounded. This is done automatically by the partitioning scheme when s is small. When s is large, a simple pigeon-hole argument works. Finally, the total number of colors has to be counted.

The results presented above leave the following problem on which there is considerable room for progress.

PROBLEM 2.5. For fixed k, close the gap between $\phi(k,n) = \Omega((\lg n/(4k))^{(k-1)}$ and $\phi(k,n) = O(n\lg^{(2k-3)}n/\lg^{(2k-4)})$.

Vishwanathan has obtained interesting results using randomized on-line algorithms.

THEOREM 2.6. *Vishwanathan[30]. There exists a randomized on-line algorithm* $\mathbb{A}$ *such that for every* k*-colorable on-line graph* $G^<$ *on* n *vertices, the expected value of* $\chi_{\mathbb{A}}(G^<) = O(k2^k n^{(k-2)/(k-1)}(lg\, n)^{1/(k-1)})$. *Moreover, for any randomized on-line algorithm* $\mathbb{B}$, *there exists a* k*-colorable on-line graph* $G^<$ *on* n *vertices such that the expected value of* $\chi_{\mathbb{B}}(G^<) = \Omega(1/(k-1)(lg\ n/(12(k+1))+1)^{k-1})$.

We comment that the algorithm $\mathbb{A}$ in Theorem 2.6 runs in polynomial time and compares well with the best off-line polynomial time approximation algorithms for graph coloring.

3. On-line χ-bounded classes.

In this section we consider classes of graphs Γ, for which there exists an on-line algorithm $\mathbb{A}$ such that $\phi_{\mathbb{A}}(k/\Gamma)$ is finite for all k. More precisely, we say that Γ is *on-line* χ*-bounded* if and only if there exists an on-line algorithm $\mathbb{A}$ and a function $g(k)$, called a χ-binding function, such that $\chi_{\mathbb{A}}(G^<) \leq g(\omega(G))$, for any on-line presentation $G^<$ of any $G \in \Gamma$. Similarly, Γ is χ*-bounded* if there exists a function $f(k)$ such that $\chi(G) \leq f(\omega(G))$, for all $G \in \Gamma$.

The results of this section have their roots in the authors' previous work in recursive combinatorics and a beautiful graph theoretical conjecture formulated independently by Gyárfás and Sumner. The problems the authors considered in recursive combinatorics can be very roughly described as follows. Given a countably infinite graph G, design an algorithm to color each vertex v of G using only certain types of local information (in particular, only finitely much information) about v. Depending on the amount of information allowed, in increasing order, the graphs may be recursive, highly recursive, or decidable. Generally, results about coloring recursive graphs, such as Bean [1], Kierstead [13], and Kierstead and Trotter[23] translate immediately to on-line results, while results on highly recursive or decidable graphs such as Kierstead[14], [15], Manaster and Rosenstein[12], and Schmerl[27] do not. The starting point for the work of this section is the following theorem. The notion of an on-line ordered set is analogous to that of an on-line graph.

THEOREM 3.1. *Kierstead[13]. There exists an on-line algorithm $\mathbb{A}$ which will partition any on-line ordered set of width at most w into $(5^w - 1)/4$ chains.*

Here we use $<$ to denote the partial order while the linear order which specifies the order in which the points are presented is L. Theorem 3.1 is proved by induction on w. The algorithm $\mathbb{A}$ first partitions the on-line ordered set $P^L = (X, <)^L$ into two parts C and X^*, where C is a maximal chain. Then an auxiliary on-line ordered set $P^{*L} = (X^*, <^*)^L$ is defined (on-line) so that the width of P^{*L} is at most $w - 1$, and each $<^*$ chain can be partitioned on-line into five $<$-chains. By induction P^L can be covered on-line by $5(5^{w-1} - 1)/4 + 1 = (5^w - 1)/4$ $<$-chains. However, the implementation of this strategy is quite complex.

Theorem 3.1 led to many new questions. One obvious problem has remained unanswered and appears to be very difficult.

PROBLEM 3.2. Is there an on-line algorithm which will partition any on-line ordered set of width at most w into $p(w)$ chains for some polynomial p?

Kierstead showed that the lower bound is non-linear and Szemerédi and Trotter showed that it is at least quadratic. Schmerl asked whether the order relation is necessary or is it the case that there is an on-line algorithm which partitions every on-line comparability graph G into a number of complete subgraphs bounded as a function of the independence number of G. This is equivalent to asking whether the class of co-comparability graphs is on-line χ-bounded. The problem is that Kierstead's chain covering algorithm makes use of the order relation between two comparable points when deciding how to "color" them. We shall see later that Schmerl's question has an affirmative answer. It is also natural to look for other classes of χ-bounded graphs and orders. Interval graphs are the co-comparability graphs of interval orders, and for this class of graphs, it is possible to obtain an exact answer.

THEOREM 3.3. *Kierstead and Trotter[23]. There is an on-line coloring algorithm which will color any on-line interval graph $G^<$ with at most $3\omega(G^<) - 2$ colors. Moreover, no on-line algorithm can do better.*

Arguing by induction on ω, one shows that $G^<$ can be partitioned on-line (just be greedy) into a maximal graph $G^{*<}$ with clique size $\omega - 1$ and a graph $H^<$ with maximum degree 2. Thus G can be colored on-line using $3(\omega - 1) - 2 + 3$ colors.

Next, we introduce the Gyárfás-Sumner Conjecture. For a graph H, let Forb(H) be the class of all graphs which do not contain an induced copy of H. Similarly,