

CSC2421: Online and Other Myopic Algorithms

Fall 2026

Allan Borodin

October 14, 2026

Lecture 9

Announcements

- I would like to agree on the composition of the groups and ask who can be more specific about an individual or group project.
- Michel He sent me a 2014 paper by S.Fung that provides a constant competitive ratio for the proportionally weighted throughput problem with revoking (and not using restarting). Michael shows that this extends to the C-Benevolent case and the case of single processing times with arbitrary weights. It should also extend to the case of D-Benevolent weights.

I will post that paper. It also contains a result where preemptions have a cost.

Today's Agenda

- Finish what I was saying about classes of graphs characterized by vertex orderings.
- Discussion of some of my recent papers with Chris Karavasilis, Michael He, Nadim Motu and David Zhang.

Finishing up what I wanted to say about graph classes characterized by bertex orderings

I was discussing open problems in Chapter 8 of Volume 2.

The first open question concerns the proportional weights interval scheduling problem on $m > 1$ machines. Perhaps surprisingly, we get better ratios for $m > 1$.

A specific question within the second open problem is whether we can obtain a constant competitive algorithm for C-Benevolent weight functions. Note: proportional weights are C-Benevolent.

The final open problem is rather vague. Given an assumption on the graph representation model, we would like to have an model for greedy algorithms when considering classes of graphs and corresponding graph problems.

For example, consider the VAM-PH and VAM-FI graph input models defined in Chapter 7 of Volume 1. We would like to re-consider graph problems (as in Volume 1, Chapter 7) in models such as the priority model and other “natural” algorithmic models we might define.

Greedy and Other Myopic Algorithms for Graph Problems

In particular, consider the VAM-FI graph input model.

The class of d -inductive graphs and the larger class of d -inductively independent graphs are defined by the existence of vertex orderings. Namely, a graph is d -inductive (respectively, d -inductively independent) if there is an ordering of the vertices $v_1, v_2, \dots, v_n$ such that $NBHD(v_i \cap \{v_{i+1}, \dots, v_n\})$ has degree at most d (respectively, has at most d independent vertices). See my 2012 paper (in my home page) with Yuli Ye called Elimination Graphs.

As an example, the class of $d = 1$ inductively independent graphs (i.e., *chordal* graphs) contains the class of interval graphs.

Once we have the appropriate ordering, we can obtain a d approximation for the unweighted Maximum Independent Set (MIS) problem. And, given the reverse of such an ordering, we can obtain a d -approximation for the graph colouring problem.

Continuing Greedy and Other Myopic Algorithms for Graph Problems

The class of d -inductive graphs includes trees with $d = 1$, planar graphs with $d = 5$, axis-oriented translates of a fixed size rectangle, and, of course, degrees of maximum degree d .

We can compute a d -inductive ordering within the VAM-FI model so that we can solve the MIS problem by adaptive priority algorithm in the VAM-FI model. Namely, we choose the next vertex v_i so that the $Nbhd(v_i)$ has minimum degree.

The class of inductively d -inductively independent graphs contains planar graphs with $d = 3$ and axis-oriented translates of rectangles of any size, graphs of tree width d and chordal graphs with $d = 1$.

We cannot compute the appropriate d -inductive ordering in the VAM-FI model. However, we could consider a more expressive input model where a vertex is represented by both the distance 1 and the distance 2 neighbourhoods of a vertex.

Finishing the discussion of Greedy and Other Myopic graph algorithm

As mentioned, the reverse of these d -inductively independent graphs provide d -approximations for the graph colouring paper.

Can we compute the reverse orderings in something like the priority model for d -inductive graphs? ?

However, we do know that we can colour d -inductive graphs in the VAM-PH model which is a more restrictive than the VAM-FI input model, and we do know how to online colour d -inductive graphs.

Beyond the graph MIS and graph colouring problems, for what other graph theoretic questions can we derive good greedy/myopic approximations? Of course, we can use a two pass-algorithm by creating the required ordering and then we have the reverse ordering.

Finally, there are other graph classes characterized (or enjoying) specified orderings of the vertices. Look up Derek Corneil's publications.

Some of my recent work

With Chris Karavasilis, we considered the unweighted interval selection problem in the online model with revoking.

This problem was essentially solved by Garay et al in their 1997 paper “Efficient On-Line Call Control Algorithms”. I say essentially as they discussed the problem in the call control model which in its generality is defined as follows:

We are given a graph with N nodes where each edge has a known capacity $c(e)$. The online requests $R_i = (s_i, t_i)$ are calls that (if accepted) are to be routed from s_i to t_i using some path P_i . Each call has a bandwidth b_i and a value or weight w_i . The objective is to maximize the weight of accepted calls subject to the constraint that for each edge e , the accepted calls do not exceed the capacity of any edge; that is, $\sum_{i: e \in P_i} b_i \leq c(e)$ for each edge e .

We discuss the general online call control problem in Volume 1, Section 4.4.

The disjoint paths problem in the line and interval selection

When $c(e) = 1$ for all edges and $b_i = 1$ for all requests, the problem is the *disjoint paths problem*.

The Garay et al paper considers the unweighted and weighted disjoint paths problem for the line network. They show that for the line network, the deterministic competitive ratio is $\log N$ where N is the number of nodes in the network. The argument for the negative result (i.e., the $\log N$ is necessary) can be restated to show that the ratio is $n/2$ where n is the length of the input sequences (i.e., the number of calls).

The interval selection problem appears to be a little more general than we can think of each interval as a call but we are not given the graph in advance. However, it is not difficult to argue that we can simulate the call admission disjoint paths on the line and the interval selection problems are essentially equivalent.

Proportional Weighted Instances

In addition to the unweighted call admission problem, Garay et al considered the proportional weights problem and showed that there is an online algorithm (with revoking) that achieves a competitive ratio of $2\phi + 1 = 2 + \sqrt{5} \approx 4.23$ and gave a lower bound of 4. Tompkins showed that $2 + \sqrt{5}$ is optimal.

Like the unweighted case, the positive and negative results for proportional weights carry over to the interval selection problem. As I believe I mentioned, it would be interesting to see if the proportional weights case can be extended to all C-Benevolent weight functions.

Following Adler and Azar, Emek, Halldórsson and Rosén gave a randomized 6-competitive algorithm for unweighted interval selection based on a 2-competitive deterministic streaming algorithm.

The Competitive Ratio in terms of the Number of Distinct Interval Lengths

We are finally ready to return to the paper by Chris Karavasilis and myself.

Lipton and Tompkins introduced the classify and randomly select method for deriving various randomized results. Specifically, they showed that in the real-time model (without revoking), for interval selection with proportional weights, there is an algorithm that achieves the ratio $O(\log \Delta)^{1+\epsilon}$ where Δ is the ratio of the longest to shortest interval.

They also showed that in the real-time model without revoking, there is a randomized 2-competitive algorithm when there are at most two distinct lengths.

Chris and I considered the unweighted online interval selection problem when there are at most k distinct interval lengths. We showed the following: The precise online (with revoking) deterministic competitive ratio, for instances with at most k distinct interval lengths, is $2k$. Note that $k \leq n$ and we already know that $n/2$ is a lower bound where n is the number of intervals in the input sequence.

The $2k$ upper and lower bounds

We first describe the algorithm that achieves the $2k$ ratio. This is arguably the simplest algorithm one can imagine and it is “memoryless”. Note also that the algorithm does not need to know k in advance.

Let I_s be the current solution. When a new interval I arrives, add I to I_s if $I \cap I_s = \emptyset$.

Otherwise, accept I if and only if $I \subset I'$ for some $I' \in I_s$ and then revoke I' .

Let OPT be an arbitrary optimal solution and let ALG represent the solution of the algorithm when it completes. We will construct a charging argument $f : OPT \rightarrow ALG$ such that every interval in ALG has at most $2k$ intervals in OPT charged to it.

Since intervals can be revoked during the algorithm's execution, we need to adaptively define the charging function f during computation. We will have *direct* and *transfer* charging. We consider the various cases when a new OPT interval I arrives.

The Charging Argument

We consider the cases when an interval $I \in OPT$ arrives.

- If $I \in ALG$, then I is mapped to itself.
- If I was rejected by the algorithm because it conflicts the ALG s current solution, we arbitrarily charge I (by direct charging) to one of the ALG intervals with which I conflicts. We can assume that I conflicts with at most two intervals in ALG s current solution $\mathcal{I}$ since otherwise it would contain some $I' \in \mathcal{I}$ and hence ALG would have taken that I' .
- If an interval I' at some time is replaced by an interval I'' subsumed by I' , then all of the current charges to I' are transferred to I'' (by what we call transfer charging).

We note that by direct charging, at most two OPT intervals can be mapped to any $I' \in \mathcal{I}$.

Finishing the charging argument and proving that $2k$ is the best ratio

We now claim that an interval in ALG can be charged by at most $2k - 2$ intervals in OPT by transfer charging. This is because at most k intervals can be involved in the nested sequence of transfer charging with the last (and smallest) interval in this sequence receiving at most $2(k - 1)$ transfer charges. This follows since at each level at most 2 intervals are added by direct charging.

We now want to show that for every the deterministic ratio is no better than $2k$. Consider a deterministic algorithm. We will show that at all times during the algorithm's execution, ALG will have at most one interval while as the depth of nesting increases OPT can add two more intervals to its solution.

Consider the following base ($k = 1$) adversarial construction. We can assume that the algorithm must take the first interval I_1 . Then the adversary inputs an interval I_2 (of the same length) on the right by a small amount.

End of lower bound argument

So now for the base case, the adversary has given two intervals with I_2 conflicting with I_1 on the right. If the algorithm rejects I_2 then the adversary inputs an interval I_3 (of the same length) conflicting I_1 on the left by a small amount. If the algorithm takes I_2 by revoking I_1 , then the adversary inputs an interval I_4 conflicting with I_2 on the right (again by a small amount).

In either case we will have ALG with one current interval I and OPT with two intervals leaving lots of room inside of I to continue this construction recursively with smaller length intervals.

It is easy to see that this argument is really a bound on the depth of nesting. With k distinct interval lengths, there can be at most $k - 1$ levels of nesting. But we could have many more distinct interval lengths with no nesting.

Finally, we can extend this lower bound argument to “memoryless” randomized algorithms. Can we eliminate the “memoryless” restriction?

Unweighted interval selection in the ROA model

Surprisingly, the same simple algorithm (as used for the adversarial sequence $2k$ ratio) has competitive ratio 2 in the random order arrival model. We conjecture that this is the optimal ratio. So far the best deterministic and randomized lower bound is $\frac{4}{3} - \epsilon$ (for any $\epsilon > 0$) in the ROA model. The lower bound is due solely to Claude.

The analysis for both the upper and lower bounds is much more substantial. It is also interesting that the same charging function can be used but, as we said, the analysis is quite different. We will want to show that $\frac{OPT}{\mathbb{E}[ALG]} \geq 2$ for all input sequences, noting that now ALG is a random variable depending on the ROA permutation whereas OPT is not effected by the random order of the sequence.

The history of the result.

The paper was rejected from a conference with one referee having suspicious about our cludgy argument. With David Zhang, we asked Claude to verify our proof and Claude showed that our argument was not correct. Claude then provided a different analysis with the improved ratio.