Interval graphs admit optimal greedy colouring in order of finish time

W. Cames van Batenburg* S. Kelk†

April 22, 2026

Abstract

This note constitutes an exercise in AI-assisted research and writing, and answers a question of the second author. It is well-known that every interval graph can be optimally properly coloured by the greedy algorithm that processes intervals in order of increasing starting time and reuses the smallest previously introduced colour whenever possible. However, the same strategy fails if instead one were to process the intervals in order of finish time. Here we show that colouring in order of finish time still produces an optimal colouring if the algorithm is given an oracle that indicates *which* previously encountered colour (not necessarily the smallest one) to reuse.

1 Introduction

Given a finite set of intervals of the real line, its *intersection graph* (or *interval graph*) has one vertex per interval and an edge between any two overlapping intervals. Interval graphs have close ties to (structural) graph theory: they are chordal graphs, perfect graphs, and they are closely related to the pathwidth of a graph, and they have been used to study scheduling problems.

For an interval $[s_i, f_i]$, we refer to s_i and f_i as its *starting time* and *finish time*, respectively. A *proper colouring* gives overlapping intervals distinct colours. It is well-known that for every interval graph, the chromatic number equals the clique number ω, the maximum number of intervals containing any single point.

A proper colouring with ω colours can be produced by a simple greedy algorithm: process the intervals in order of *increasing starting time* and, for each interval in turn, reuse the *smallest* indexed colour not already used by an overlapping earlier interval, or introduce a new one if no such colour exists. When it's the turn of interval I to be coloured, all previously coloured intervals that intersect I must be mutually intersecting. This means that at most ω-1 previously encountered intervals intersect I, so that a free colour is always available for I.

However, this greedy strategy fails if one were to colour in order of *increasing finish time*, while still choosing the smallest available colour at each step. In that case the previously encountered intervals intersecting I do not need to be mutually overlapping, so there might be more than $\omega - 1$ of them. Here is a small example with $\omega = 2$ where this strategy indeed needs more than ω colours.

Example 1.1. Consider the bipartite interval graph with intervals

$$A = [0,1], \quad B = [\tfrac{1}{2}, 2], \quad D = [\tfrac{5}{2}, 3], \quad C = [\tfrac{3}{2}, \tfrac{7}{2}],$$

*Département d'Informatique, Université libre de Bruxelles, Belgium (w.p.s.camesvanbatenburg@gmail.com).
Supported by the Belgian National Fund for Scientific Research (FNRS).
†Department of Advanced Computing Sciences, Maastricht University, The Netherlands (steven.kelk@maastrichtuniversity.nl).

processed in their finish time order A, B, D, C. The algorithm assigns $A \mapsto 1$ and $B \mapsto 2$ (since B overlaps A). Then D overlaps neither A nor B, so the smallest available colour is 1. Finally C overlaps both B (colour 2) and D (colour 1), so this forces the introduction of a third colour.

The question is whether this kind of mishap can always be avoided by a cleverer choice. In a post on the social medium Mathstodon [1], the second author asked if optimal colouring is possible in the following situation. Define *non-deterministic greedy by finish time* as the algorithm above except that, whenever at least one already-introduced colour is available for the current interval, an oracle decides which of the available colours to reuse; the algorithm is still forbidden from introducing a new colour if any old one is available, but it no longer needs to choose the *smallest* among the old available colours. That algorithm does produce an optimal colouring.

Theorem 1.2. *For any finite set of intervals whose intersection graph has clique number ω, non-deterministic greedy colouring by finish time can colour them properly using at most ω colours.*

The proof fixes an arbitrary optimal colouring φ and builds an execution γ of non-deterministic greedy, maintaining a partial injection σ from the colours that γ has introduced to the colours of φ. So $\sigma(c)$ records the φ-colour of the most recent interval that received γ-colour c. When a new interval is processed, the algorithm prefers to assign it a γ-colour c with $\sigma(c) = \varphi(I)$, which a short lemma shows is always available in that case; only when this is impossible is a new γ-colour introduced. Since σ is an injection into $\{1, \ldots, \omega\}$, at most ω colours are ever introduced.

AI assistance
This note was written with help of the large language model Claude Opus 4.6. The authors introduced the outlines of the proof idea. The model then formulated precise induction hypotheses and drafted the argument and write-up. The majority of the necessary rewriting has been done via prompts rather than direct human modification. An autoformalization of the proof in Lean 4, produced with help of *Aristotle (Harmonic)*, is available in the accompanying repository [2].

We are not aware of any application and so, at least for now, we view the result itself as mostly a combinatorial curiosity: we do not claim any deep consequence or novelty. In particular, there is a nonzero probability that the same argument already existed somewhere in the training data. Still, it seemed worth recording, both because the question of the second author had a cleanly positive answer and because it provided a useful, verifiable test bed for AI-assisted mathematical writing and research.

2 The proof

Proof of Theorem 1.2. Since interval graphs are perfect, fix a proper colouring $\varphi\colon \{I_1, \ldots, I_n\} \to \{1, \ldots, \omega\}$, where $I_1, \ldots, I_n$ are ordered by increasing finish time. We construct a colouring $\gamma\colon \{I_1, \ldots, I_n\} \to \{1, \ldots, \omega\}$ inductively and show it is a valid output of non-deterministic greedy.

We maintain a set $S \subseteq \{1, \ldots, \omega\}$ of *introduced* colours and an injection $\sigma\colon S \to \{1, \ldots, \omega\}$. A colour $c \in S$ is *available* for an interval I if no interval among those already coloured with γ-value c overlaps I.

Claim 2.1. *Suppose γ has been defined on $I_1, \ldots, I_k$, is proper, and $c \in S$. Let J be the last of $I_1, \ldots, I_k$ with $\gamma(J) = c$. If a future interval I does not overlap J, then I overlaps no interval among $I_1, \ldots, I_k$ with γ-value c.*

Proof. Let $J_1, \ldots, J_m = J$ be the intervals with γ-value c, in processing order. Since γ is proper, $\text{start}(J_{\ell+1}) \geq \text{finish}(J_\ell)$ for each ℓ. Suppose I overlaps some J_ℓ for $\ell < m$. Then

$$\text{start}(I) < \text{finish}(J_\ell) \leq \text{start}(J_m) < \text{finish}(J_m) \leq \text{finish}(I),$$

so I overlaps $J_m = J$, a contradiction. ♦

We prove by induction on k that after defining γ on $I_1, \ldots, I_k$, the following conditions hold:

(H1) γ is proper.

(H2) The greedy rule has been respected: for each $\ell \leq k$, either $\gamma(I_\ell) \in S$ was available for I_ℓ at the time I_ℓ was processed, or $\gamma(I_\ell)$ was a newly introduced colour and no colour in S was available.

(H3) $\sigma \colon S \to \{1, \ldots, \omega\}$ is injective.

(H4) For each $c \in S$, the last interval J among $I_1, \ldots, I_k$ with $\gamma(J) = c$ satisfies $\varphi(J) = \sigma(c)$.

Base case. When $k = 0$, we have $S = \varnothing$ and σ is the empty function. All conditions hold vacuously.

Induction step. Assume (H1)–(H4) hold after processing $I_1, \ldots, I_k$. Let $j = \varphi(I_{k+1})$.

Case A: $j \in \text{im}(\sigma)$. Let $c = \sigma^{-1}(j)$ and let J be the last interval with $\gamma(J) = c$. By (H4), $\varphi(J) = \sigma(c) = j = \varphi(I_{k+1})$. Since φ is proper, J and I_{k+1} do not overlap. By Claim 2.1, c is available for I_{k+1}. Define $\gamma(I_{k+1}) = c$. Since $c \in S$ is available, (H2) is satisfied. Since I_{k+1} overlaps no interval with γ-value c, (H1) is preserved. Since $\varphi(I_{k+1}) = j = \sigma(c)$, (H4) is preserved. Neither S nor σ changes, so (H3) is preserved.

Case B: $j \notin \text{im}(\sigma)$ *and some colour in* S *is available for* I_{k+1}. Let $c' \in S$ be available for I_{k+1}. Define $\gamma(I_{k+1}) = c'$. Since $c' \in S$ is available, (H2) is satisfied. Since c' is available, (H1) is preserved. Update σ by redefining $\sigma(c') = j$; since $j \notin \text{im}(\sigma)$ before the update and only the value at c' changes, σ remains injective, preserving (H3). The last interval with γ-value c' is now I_{k+1}, and $\varphi(I_{k+1}) = j = \sigma(c')$, preserving (H4).

Case C: $j \notin \text{im}(\sigma)$ *and no colour in* S *is available for* I_{k+1}. Since σ is injective and $j \notin \text{im}(\sigma) \subseteq \{1, \ldots, \omega\}$, we have $|S| = |\text{im}(\sigma)| \leq \omega - 1$, so $\{1, \ldots, \omega\} \setminus S$ is nonempty. Choose $c_{\text{new}} \in \{1, \ldots, \omega\} \setminus S$. Define $\gamma(I_{k+1}) = c_{\text{new}}$ and add c_{new} to S. Since no colour in S was available, introducing c_{new} respects (H2). Since c_{new} was not previously used, (H1) is preserved. Set $\sigma(c_{\text{new}}) = j$; since c_{new} is new to S and $j \notin \text{im}(\sigma)$, the extended σ remains injective, preserving (H3). The last interval with γ-value c_{new} is I_{k+1}, and $\varphi(I_{k+1}) = j = \sigma(c_{\text{new}})$, preserving (H4).

In all cases, $S \subseteq \{1, \ldots, \omega\}$, so γ takes values in $\{1, \ldots, \omega\}$ and uses at most ω colours. □

References

[1] S. Kelk, Mathstodon post, `https://mathstodon.xyz/@skelk/116174927025037600`, March 2026.

[2] Companion repository with a Lean 4 formalization of Theorem 1.2, `https://woutercvb.github.io/greedy_by_finish_time.txt`, 2026.