Monotone Paths in Dense Edge-Ordered Graphs

Kevin G. Milans *

October 13, 2015

Abstract

The *altitude* of a graph G, denoted $f(G)$, is the largest integer k such that under each ordering of $E(G)$, there exists a path of length k which traverses edges in increasing order. In 1971, Chvátal and Komlós asked for $f(K_n)$, where K_n is the complete graph on n vertices. In 1973, Graham and Kleitman proved that $f(K_n) \geq \sqrt{n-3/4} - 1/2$ and in 1984, Calderbank, Chung, and Sturtevant proved that $f(K_n) \leq (\frac{1}{2}+o(1))n$. We show that $f(K_n) \geq (\frac{1}{20}-o(1))(n/\lg n)^{2/3}$.

1 Introduction

A *totally ordered graph* is a graph G that is associated with a total ordering of its vertex set $V(G)$ and a total ordering of its edge set $E(G)$. We use $T(G)$ and $T'(G)$ to denote the total orderings of $V(G)$ and $E(G)$ respectively. When only the vertices or only the edges of G are totally ordered, we call G an *ordered graph* or an *edge-ordered graph*, respectively. An *ordering*, *edge-ordering*, or *total ordering* of a graph G is an ordered, edge-ordered, or totally ordered graph whose underlying graph is G.

In an edge-ordered graph G, a *monotone path* is a path which traverses edges in increasing order with respect to $T'(G)$. A *monotone trail* is similar, except that a trail is allowed to revisit vertices. The *altitude* of a graph G, denoted $f(G)$, is the maximum integer k such that every edge-ordering of G contains a monotone path of length k. Also, let $f^\star(G)$ be the maximum integer k such that every edge-ordering of G contains a monotone trail of length k.

In 1971, Chvátal and Komlós [4] asked for $f(K_n)$ and $f^\star(K_n)$, where K_n denotes the complete graph on n vertices. Citing private communication, Chvátal and Komlós noted in their 1971 paper that Graham and Kleitman had already proved $\Omega(n^{1/2}) \leq f(K_n) < (\frac{3}{4}+\varepsilon)n$ and established $f^\star(K_n)$ exactly: $f^\star(K_n) = n-1$ unless $n \in \{3,5\}$, in which case $f^\star(K_n) = n$.

To show $f^\star(K_n) \geq n-1$, Graham and Kleitman [6] proved that if G has average degree d, then $f^\star(G) \geq d$. Friedgut communicated to Winkler [13] an elegant formulation of their proof, known as the *pedestrian argument*. For an n-vertex edge-ordered graph G, the pedestrian argument involves n pedestrians, with one starting at each vertex in G. An announcer calls out the names of the edges in order according to $T'(G)$. When e is called, both pedestrians at the endpoints of e traverse e, trading places. Since each pedestrian travels along a monotone trail and each edge is traversed by two pedestrians, the average length of a pedestrian's monotone trail is $2|E(G)|/n$, which equals d. The pedestrian argument has recently been modified to produce monotone paths (see [8] and [5]).

*West Virginia University, Morgantown, WV 26505, (milans@math.wvu.edu).

Determining the altitude of a graph appears to be difficult in general. In 1973, Graham and Kleitman [6] published their results on $f(K_n)$ and $f^\star(G)$. In particular, they proved that $\sqrt{n-3/4}-1/2 \le f(K_n) < 3n/4$, and they conjectured that $f(K_n)$ is closer to their upper bound than their lower bound. They also commented that, with additional effort, their lower bound could be improved to $f(K_n) \ge (c-o(1))\sqrt{n}$ for some $c > 1$. In his Master's thesis from the same year, Rödl [11] proved that if G has average degree d, then $f(G) \ge (1-o(1))\sqrt{d}$; for $G = K_n$, Rödl's result matches the Graham–Kleitman lower bound asymptotically. Rödl also noticed that the ideas in the Graham–Kleitman upper bound can be combined with results in design theory to prove $f(K_n) \le (\frac{2}{3}+o(1))n$. Alspach, Heinrich, and Graham (unpublished, see [3]) further improved the upper bound to $f(G) \le (\frac{7}{12}+o(1))n$. In 1984, Calderbank, Chung, and Sturtevant [3] obtained the best known upper bound: $f(K_n) \le (\frac{1}{2}+o(1))n$. After 1984, explicit progress on determining $f(K_n)$ slowed (but see [2] for exact values for $n \le 8$). In the meantime, other interesting results on the altitude of graphs have appeared.

In 2001, Roditty, Shoham, and Yuster [10] proved that $f(G) \le 9$ if G is planar and showed that $f(C_n \vee \overline{K_2}) \ge 5$ for $n \ge 99$, where $C_n \vee \overline{K_2}$ is the planar graph obtained by joining the n-vertex cycle C_n and a pair of non-adjacent vertices. Consequently, the maximum of altitude of a planar graph is between 5 and 9.

Clearly, $f(G) \le f^\star(G)$. The *edge-chromatic number* of G, denoted $\chi'(G)$, is the minimum k such that $E(G)$ is the union of k matchings. Ordering $E(G)$ so that each matching is an interval shows that $f^\star(G) \le \chi'(G)$. Vizing's theorem [12] states that $\chi'(G) \le \Delta(G)+1$, where $\Delta(G)$ is the maximum degree of G. It follows that $f(G) \le \Delta(G)+1$.

Improving a result of Yuster [14], Alon [1] gave a short proof that there exist k-regular graphs G with $f(G) \ge k$, as follows. The *girth* of G is the length of a shortest cycle in G. If G has girth g, then every trail of length less than g is a path. Therefore $f(G) \ge \min\{g-1, f^\star(G)\} \ge \min\{g-1, d\}$, where d is the average degree of G. In particular, if G is k-regular and has girth larger than k, then $f(G) \ge k$. For $k = 3$, better constructions are known. Mynhardt, Burger, Clark, Falvai, and Henderson [9] characterized the 3-regular graphs with girth at least 5 and altitude 3, and then used the characterization to show that the flower snarks are examples of 3-regular graphs with altitude 4. For $k \ge 4$, it remains open to decide whether there are graphs G with $\Delta(G) = k$ and $f(G) = k+1$.

A *Hamiltonian path* in a graph is a path containing all of its vertices. Katrenič and and Semanišin [7] proved that deciding whether a given edge-ordered graph contains a Hamiltonian monotone path is NP-complete. Although it seems likely that computing the altitude of a given graph is NP-hard or worse, we note that the result of Katrenič and Semanišin does not directly imply this.

Lavrov and Loh [8] investigated the maximum length of a monotone path in a random edge-ordering of K_n. They showed that with probability tending to 1, a random edge-ordering of K_n contains a monotone path of length at least $0.85n$. Consequently, edge-orderings of K_n that give sublinear upper bounds on $f(K_n)$, if they exist, are rare. They also proved that with probability at least $1/e - o(1)$, a random edge-ordering of K_n contains a Hamiltonian monotone path. The common strengthening of these results leads to a natural and beautiful conjecture.

Conjecture 1.1 (Lavrov–Loh [8])**.** With probability tending to 1, a random edge-ordering of K_n contains a Hamiltonian monotone path.

Recently, De Silva, Molla, Pfender, Retter, and Tait [5] proved that $f(Q_n) \ge n/\lg n$ where Q_n is the n-dimensional hypercube and lg denotes the base-2 logarithm. They also showed that if

$\omega(n) \to \infty$ and $p \le (\omega(n) \ln n)/n^{1/2}$, then with probability tending to 1 the Erdős–Rényi random graph $G(n,p)$ has altitude at least $(1-o(1))\frac{np}{\omega(n)\ln n}$. Consequently, there are graphs with average degree $\sqrt{n}(\ln n)^2$ and altitude at least $(1-o(1))\sqrt{n}$. These graphs are sparse and yet the lower bound on their altitude asymptotically matches the lower bound on $f(K_n)$ due to Graham and Kleitman.

In this paper, we improve Rödl's result for sufficiently dense graphs. We show that if G is an n-vertex graph with average degree d and $s^2/d \to 0$ where $s = \Theta(n^{1/3}(\log n)^{2/3})$, then $f(G) \ge (1-o(1))\frac{d}{4s}$. For $G = K_n$, we obtain $f(K_n) \ge (\frac{1}{20} - o(1))(n/\lg n)^{2/3}$. Our proof is based on a simple algorithm to extend monotone paths.

2 Monotone Path Algorithm

In his Master's thesis, Rödl [11] gave an elegant argument that $f(G) \ge (1-o(1))\sqrt{d}$ where d is the average degree of G, which we outline as follows. Let G be an edge-ordered graph with average degree d and suppose that k is an integer with $d \ge 2\binom{k+1}{2} = 2(1+\cdots+k)$. Obtain G' from G by marking at each vertex v the k largest edges incident to v (or all edges incident to v if $d(v) < k$) and then removing all marked edges. Since G' has average degree at least $d-2k$, by induction G' contains a monotone path $x_0 \ldots x_{k-1}$ of length $k-1$. Since x_{k-1} is not isolated in G', it follows that x_{k-1} is incident to at least k edges in $E(G) - E(G')$, and one of these extends $x_0 \ldots x_{k-1}$ to a monotone path of length k. Rödl's idea of reserving large edges at each vertex for path extension plays a key role in our approach. We make a slight change in that we require the vertices to have disjoint sets of reserved edges. We organize the edges in a table.

Let G be a totally ordered graph. The *height table* of G is an array A whose columns are indexed by $V(G)$ and rows are indexed by the positive integers. Each cell in A is empty or contains an edge in G. For $u \in V(G)$ and a positive integer i, we use $A(i,u)$ to denote the contents of the cell in A located in row i and column u. We order the cells of A so that $A(i,u)$ precedes $A(i',u')$ if and only if $i < i'$ or $i = i'$ and u precedes u' in $T(G)$. We define A iteratively. Given that the contents of all preceding cells have been defined, let $A(i,u)$ be the largest edge (relative to $T'(G)$) incident to u not appearing in a preceding cell; if no such edge exists, then $A(i,u)$ is empty. Note that each edge appears in exactly one cell in A. We define the *height* of e in G, denoted $h_G(e)$, to be the index of the row in A containing e.

Extending a given monotone path is a key step in our algorithm. The *height* of a nontrivial monotone path $x_0 \ldots x_k$ is the height of its last edge $x_{k-1}x_k$.

Lemma 2.1. *Let G be a totally ordered graph. For $1 \le k < r$, each monotone path of length k and height r extends to a monotone path of length $k+1$ and height at least $r-k$.*

Proof. Let A be the height table of G, and let $x_0 \ldots x_k$ be a monotone path of length k and height r. Since $x_{k-1}x_k$ appears in row r in A, this edge did not already appear when $A(i,x_k)$ is defined for $i < r$. It follows that for $i < r$, the cell $A(i,x_k)$ contains an edge incident to x_k which is larger than $x_{k-1}x_k$ in $T'(G)$. Let $S = \{A(i,x_k)\colon\ r-k \le i \le r-1\}$. Since $|S| = k$ and $x_{k-1}x_k \notin S$, some edge in S joins x_k with a vertex outside $\{x_0,\ldots,x_{k-1}\}$ and extends the path as claimed. □

Starting with a single edge and iterating Lemma 2.1, we obtain the following.

Lemma 2.2. *Let G be a totally ordered graph and let x_0x_1 be an edge in G of height r. If t is a positive integer and $\binom{t}{2} < r$, then G contains monotone path $x_0x_1\ldots x_t$ of height at least $r - \binom{t}{2}$.*

Proof. By induction on t. The lemma is clear when $t = 1$. For $t > 1$, the inductive hypothesis implies that G contains a monotone path $x_0 x_1 \dots x_{t-1}$ of height at least $r - \binom{t-1}{2}$. With $k = t - 1$, we apply Lemma 2.1 to obtain a monotone path $x_0 \dots x_t$ with height at least $(r - \binom{t-1}{2}) - (t-1)$ which equals $r - \binom{t}{2}$. □

Using Lemma 2.2, we match Rödl's bound $f(G) \geq (1 - o(1))\sqrt{d}$ asymptotically. We include the short proof for completeness.

Theorem 2.3. *If G has average degree d, then $f(G) \geq \lfloor 1/2 + \sqrt{d} \rfloor$.*

Proof. Let H be a total ordering of G, and let $x_0 x_1$ be an edge of maximum height r. Since each row of the height table contains n cells, it follows that $r \geq |E(G)|/n = d/2$. If t is a positive integer and $\binom{t}{2} < d/2$, then we may apply Lemma 2.2 to extend $x_0 x_1$ to a monotone path of length t in H. Hence, $\binom{t}{2} < d/2$ implies that $f(G) \geq t$. With $t = \lfloor 1/2 + \sqrt{d} \rfloor$, we have that $\binom{t}{2} < d/2$ and therefore $f(G) \geq \lfloor 1/2 + \sqrt{d} \rfloor$. □

Let G be a totally ordered graph and let $x_0 \dots x_k$ be a monotone path in G. Viewing height as a resource, extending $x_0 \dots x_k$ becomes more expensive as k grows. When extending becomes too expensive, we delete $\{x_0, \dots, x_{k-2}\}$ from G to form a new totally ordered graph G' (which inherits the orderings of $V(G)$ and $E(G)$), and we extend $x_{k-1} x_k$ to a monotone path in G'. For this to work, we must show that the height of $x_{k-1} x_k$ does not decrease too much when we delete $\{x_0, \dots, x_{k-2}\}$ from G.

Definition 2.4. Let G be a totally ordered graph. For $S \subseteq V(G)$ and an edge e in $G - S$, we define $\mathrm{drop}(G, S, e)$ to be $h_G(e) - h_{G-S}(e)$. For $s \leq n - 2$, let $g(n, s)$ be the maximum of $\mathrm{drop}(G, S, e)$ over all n-vertex totally ordered graphs G, all sets S of s vertices in G, and all edges $e \in E(G - S)$.

Note that $g(n, s)$ is monotonic in n, since adding isolated vertices to a totally ordered graph G and inserting them arbitrarily into the vertex ordering gives a larger totally ordered graph G' such that $\mathrm{drop}(G, S, e) = \mathrm{drop}(G', S, e)$ for all $S \subseteq V(G)$ and $e \in E(G - S)$.

Lemma 2.5. *Let G be an n-vertex totally ordered graph and let $x_0 x_1$ be an edge of height r. If s is a positive integer and $s \leq n - 2$, then G contains a monotone path extending $x_0 x_1$ of length at least $sk + 1$, where $k = \lfloor (r-1)/(\binom{s+1}{2} + g(n, s)) \rfloor$.*

Proof. By induction on n. If $k = 0$, then the lemma is clear. Otherwise, $r - 1 \geq \binom{s+1}{2} + g(n, s)$ and we may apply Lemma 2.2 to obtain a monotone path $x_0 \dots x_{s+1}$ of height at least $r - \binom{s+1}{2}$. Let $S = \{x_0, \dots, x_{s-1}\}$ and let $G' = G - S$. We have that $h_{G'}(x_s x_{s+1}) = h_G(x_s x_{s+1}) - \mathrm{drop}(G, S, x_s x_{s+1}) \geq r - \binom{s+1}{2} - g(n, s)$.

Applying the inductive hypothesis to G' and $x_s x_{s+1}$, we obtain a monotone path P' in G' extending $x_s x_{s+1}$ of length at least $sk' + 1$, where

$$k' = \left\lfloor \frac{r - \binom{s+1}{2} - g(n,s) - 1}{\binom{s+1}{2} + g(n-s, s)} \right\rfloor \geq \left\lfloor \frac{r - \binom{s+1}{2} - g(n,s) - 1}{\binom{s+1}{2} + g(n, s)} \right\rfloor = k - 1.$$

Prepending $x_0 \dots x_s$ to P' produces a monotone path in G of length at least $s + sk' + 1$, and $s + sk' + 1 \geq sk + 1$. □

3 The Token Game

Our goal is to prove an upper bound on $g(n,s)$. Let G be an n-vertex totally ordered graph, and let S be a set of s vertices of G. We analyze an iterative process which obtains the height table of $G-S$ from the height table of G. Let $G' = G - S$, let A be the array obtained from the height table of G by deleting columns indexed by vertices in S, and let A' be the height table of G'. Note that the cells of both A and A' are indexed by Z, where $Z = \{1,2,3,\ldots\} \times V(G')$. We order Z in the same order as the corresponding cells in A' are defined; that is, $(i,u) \le (i',v)$ if and only if $i < i'$ or $i = i'$ and $u \le v$ in $T(G')$. For $\beta \in Z$, the *open down-set* of β, denoted $D(\beta)$, is $\{\gamma \in Z\colon\ \gamma < \beta\}$ and the *closed up-set* of β, denoted $U[\beta]$, is $\{\gamma \in Z\colon\ \gamma \ge \beta\}$. Similarly, the *interval* $[\beta,\gamma]$ is $\{\delta \in Z\colon\ \beta \le \delta \le \gamma\}$.

We produce a sequence of arrays $\{A_\beta\colon\ \beta \in Z\}$ which initially resemble A and later resemble A'. For $\beta \in Z$, the cells of A_β are indexed by Z and are partitioned into a *lower part* indexed by $D(\beta)$ and an *upper part* indexed by $U[\beta]$.

For $\beta \in Z$, each cell in A_β is either empty, contains an edge in G', or contains an object called a *hole*. Moreover, each edge in G' appears in one cell in A_β. Each A_β also has a *critical interval* $[(i,u),(j,u)]$, where $\beta = (i,u)$ and j is the least integer such that $j \ge i$ and $A_\beta(j,u)$ does not contain a hole.

Lemma 3.1. *There is a sequence of arrays $\{A_\beta\colon\ \beta \in Z\}$ such that each column in the initial array has at most s holes, and for each $\beta \in Z$ the following hold.*

1. *If $\delta < \beta$, then $A_\beta(\delta) = A'(\delta)$.*
2. *If $\delta \ge \beta$ and $A_\beta(\delta)$ does not contain a hole, then $A_\beta(\delta) = A(\delta)$.*
3. *If γ is the successor of β in Z, then A_γ is obtained from A_β by swapping $A_\beta(\beta)$ and $A_\beta(\delta)$, where δ is in the critical interval of A_β. Moreover, if β and δ index cells in distinct columns u and v, then $A_\beta(\delta) = uv$.*

Proof. Recall that A is obtained from the height table of G by deleting columns indexed by vertices in S. Note that A omits every edge with both endpoints in S and contains every edge in G'. An edge $uv \in [S,\overline{S}]$ with $u \notin S$ and $v \in S$ appears in A if and only if uv is in column u in the height table of G. Let α be the minimum element in Z, and let A_α be the array obtained from A by replacing edges in $[S,\overline{S}]$ with holes. If u indexes a column in A_α, then each hole in column u replaces an edge uv in G with $v \in S$, and therefore each column in A_α contains at most s holes. Clearly, every edge in G' appears once in A_α and A_α satisfies properties (1) and (2).

We obtain other arrays iteratively. Let $\beta = (i,u)$, let γ be the successor of β, and suppose that A_β has been previously defined but A_γ is not yet defined. Analogously to A_β, we partition of the cells of A' into a lower part indexed by $D(\beta)$ and an upper part indexed by $U[\beta]$. Since A_β and A' contain the same set of edges and agree on their lower parts, it follows that the upper parts of A_β and A' contain the same edges (possibly in a different order). We consider two cases, depending on whether $A'(\beta)$ is empty or contains an edge in G'.

Case 1: $A'(\beta)$ is not empty. Let $e = A'(\beta)$, and let δ be the index of the cell in A_β containing e. We claim that δ is in the critical interval $[(i,u),(j,u)]$ of A_β. Since e is in the upper part of A', it follows that e is in the upper part of A_β and so $\delta \ge \beta = (i,u)$. Since $\delta,(j,u) \in U[\beta]$ and neither $A_\beta(\delta)$ nor $A_\beta(j,u)$ contains a hole, it follows from (2) that $A(\delta) = A_\beta(\delta) = e$ and $A(j,u) = A_\beta(j,u)$. Suppose for a contradiction that $\delta > (j,u)$. Note that e is available for $A(j,u)$

when building the height table of G, and so $A(j,u) = e'$ for some edge e' incident to u such that $e' > e$ in $T'(G)$. Since $A_\beta(j,u) = A(j,u) = e'$, it follows that both e and e' appear in the upper part of A_β and hence in the upper part of A' also. Therefore both e and e' are available for $A'(\beta)$ when building the height table of G'. The selection of e over e' for $A'(\beta)$ implies that $e > e'$ in $T'(G')$, contradicting that $e' > e$ in $T'(G)$. Therefore $\delta \le (j,u)$ and δ is in the critical interval of A_β as claimed. Obtain A_γ from A_β by swapping the contents of cells $A_\beta(\beta)$ and $A_\beta(\delta)$ (if $\beta = \delta$, then $A_\gamma = A_\beta$). Note that if δ indexes a cell in column v and $v \ne u$, then $A'(\beta) = e$ and $A(\delta) = e$ imply that e is incident to both u and v, so that $A_\beta(\delta) = e = uv$, satisfying (3).

We check that A_γ satisfies (1) and (2). Since γ is the successor of β and $A_\gamma(\beta) = A_\beta(\delta) = e = A'(\beta)$, it follows that A_γ satisfies (1). If the critical interval $[(i,u),(j,u)]$ of A_β has size 1, then $\beta = (i,u) = \delta = (j,u)$ and $A_\gamma = A_\beta$, implying that A_γ satisfies (2). Otherwise $j > i$ and $A_\beta(\beta)$ contains a hole. Relative to A_β, the only change in the upper part of A_γ is that $A_\gamma(\delta)$ becomes a hole after swapping $A_\beta(\beta)$ and $A_\beta(\delta)$, and so A_γ satisfies (2).

Case 2: $A'(\beta)$ is empty. This implies that the upper part of A' contains no edge incident to u, and so the upper part of A_β also contains no edge incident to u. In particular, $A_\beta(j,u)$ is empty, where $[(i,u),(j,u)]$ is the critical interval of A_β. We obtain A_γ from A_β by swapping the contents of cells $A_\beta(i,u)$ and $A_\beta(j,u)$, satisfying (3). Since $A_\gamma(\beta)$ and $A'(\beta)$ are both empty, A_γ satisfies (1). Relative to A_β, the upper part of A_γ is either unchanged or contains a new hole at $A_\gamma(j,u)$. It follows that A_γ also satisfies (2). □

Given the sequence of arrays $\{A_\beta \colon \beta \in Z\}$ from Lemma 3.1, we obtain a useful upper bound on $\mathrm{drop}(G,S,e)$.

Lemma 3.2. *Let e be an edge in G' and choose $\beta \in Z$ so that $A'(\beta) = e$. If $[(i,u),(j,u)]$ is the critical interval of A_β, then* $\mathrm{drop}(G,S,e) \le j - i$.

Proof. Since $\beta = (i,u)$ and e appears in row i of the height table of G', it follows that $h_{G'}(e) = i$. Let δ index the cell in A_β containing e. Since the successor A_γ of A_β satisfies $A_\gamma(\beta) = A'(\beta) = e$, it follows that A_γ is obtained from A_β by swapping $A_\beta(\beta)$ with $A_\beta(\delta)$. By (3), we have that δ is in the critical interval $[(i,u),(j,u)]$ of A_β, and so $\delta = (\ell, v)$ where $i \le \ell \le j$. Since $\delta \ge \beta$ and $A_\beta(\delta)$ is not a hole, by (2) we have that $e = A_\beta(\delta) = A(\delta)$. Therefore e appears in row ℓ of the height table of G and so $h_G(e) = \ell$. We conclude $\mathrm{drop}(G,S,e) = \ell - i \le j - i$. □

We define the *height* of a critical interval $[(i,u),(j,u)]$ to be $j - i$. Note that the height of the critical interval of A_β is at most the number of holes in column u of A_β. Also, by property (1) of Lemma 3.1, all holes of A_β are contained in the upper part of A_β. Analyzing the movement of the holes as β increases in Z naturally leads to a single player game.

A *token game* is a game played on an array B with rows indexed by the positive integers and columns indexed by a finite list. Let $B(i,u)$ denote the cell in row i and column u. Each cell in B is empty or contains a *token*. A token in cell $B(i,u)$ is *grounded* if all cells in column u below $B(i,u)$ contain tokens; a token which is not grounded is *ungrounded*. One of the columns is distinguished as the *active column*.

A step in a token game modifies B to produce a new array B', subject to certain rules. Let u be the active column. If column u contains grounded tokens, then the player may optionally move the highest grounded token in column u from its cell $B(i,u)$ to an empty cell $B(i',v)$, provided that $i' \le i$ and no prior step in the game moved a token between columns u and v. Next, all ungrounded tokens in column u shift down by one cell, and the active column advances cyclically. A step in which a token moves between columns is a *transfer step*. The list of arrays produced in a token game is its *transcript*.

An (n, s)-*token game* is a token game with n columns, each of which initially contains at most s tokens. Let $\hat{g}(n, s)$ be the maximum number of tokens that can be placed in a single column in an (n, s)-token game. The following gives the connection between $g(n, s)$ and $\hat{g}(n, s)$.

Lemma 3.3. $g(n, s) \leq \hat{g}(n - s, s)$

Proof. Let G be an n-vertex totally ordered graph and let S be a set of s vertices in G such that $\text{drop}(G, S, e) = g(n, s)$ for some edge e in $G - S$. Let $G' = G - S$, let A' be the height table of G', obtain A from the height table of G by deleting columns indexed by S, and apply Lemma 3.1 to obtain the sequence of arrays $\{A_\beta \colon \beta \in Z\}$. We use this sequence to play the $(n - s, s)$-token game so that at least $g(n, s)$ tokens are placed in some column.

Construct a sequence $\{B_\beta \colon \beta \in Z\}$ of token arrays as follows. Let $\beta = (i, u)$. We put a token in $B_\beta(j, v)$ if and only if $A_\beta(k, v)$ contains a hole, where $k = j + i$ if $v < u$ in $T(G')$ and $k = j + i - 1$ otherwise. Equivalently, we obtain B_β from A_β by removing all edges so that only holes and empty cells remain, shifting cells down to discard the lower part of A_β, and replacing holes with tokens.

We claim that the sequence $\{B_\beta \colon \beta \in Z\}$ is the transcript of an $(n - s, s)$-token game in which the active column of B_β is the second coordinate in β. Let α be the minimum element in Z, and note that each column in A_α contains at most s holes by Lemma 3.1. It follows that each column in B_α contains at most s tokens, satisfying the initial condition of an $(n - s, s)$-token game.

Let $\beta = (i, u)$ and let γ be the successor of β. From property (3) of Lemma 3.1, we have that A_γ is obtained from A_β by swapping $A_\beta(\beta)$ and $A_\beta(\delta)$ for some δ in the critical interval $[(i, u), (j, u)]$ of A_β. If the critical interval has size 1, then $A_\gamma = A_\beta$ and column u of B_β contains no grounded tokens. We obtain B_γ from B_β by allowing the tokens in column u to shift down by 1 cell. The active column advances, completing a legal move in the token game.

Otherwise $j > i$. Recall that the cells of B_β correspond to the upper part of A_β. The cells indexed by the critical interval $[(i, u), (j, u)]$ of A_β correspond to the cells in B_β of height at most $j - i$, except that the last cell $A_\beta(j, u)$ corresponds to $B_\beta(j - i + 1, u)$ which has height $j - i + 1$.

Since $A_\beta(\ell, u)$ contains a hole for $i \leq \ell < j$, it follows that $B_\beta(\ell, u)$ contains a grounded token for $1 \leq \ell \leq j - i$. Since $A_\beta(\beta)$ contains a hole and $A_\beta(\delta)$ does not, it follows that $\delta > \beta$ and we obtain A_γ from A_β by swapping the contents of distinct cells $A_\beta(\beta)$ and $A_\beta(\delta)$. Therefore we obtain B_γ from B_β by firstly moving the grounded token in $B_\beta(1, u)$ to an empty cell of height at most $j - i$ or to $B_\beta(j - i + 1, u)$ and secondly shifting the contents of all cells in column u down by 1 cell. Equivalently, we obtain B_γ from B_β by optionally moving the highest grounded token from $B_\beta(j - i, u)$ to an empty cell of height at most $j - i$ and shifting the ungrounded tokens in column u down by 1 cell. This is allowed in a token game provided that we have not executed a transfer step between a pair of columns more than once.

Suppose that the transition from B_β to B_γ represents the first transfer step between distinct columns u and v; we may assume without loss of generality that a token is moved from column u in B_β to column v in B_γ. It follows that a hole in $A_\beta(\beta)$ is swapped with the contents of $A_\beta(\delta)$ to form A_γ, where β and δ index cells in columns u and v respectively. By property (3) of Lemma 3.1, we have that $A_\beta(\delta) = uv$. Since $\delta \geq \beta$, the edge uv is in the upper part of A_β. On the other hand, we have $\beta < \gamma$ and $A_\gamma(\beta) = A_\beta(\delta) = uv$, and so uv is in the lower part of A_γ. In fact, $A_{\gamma'}(\beta) = A_\gamma(\beta) = uv$ for $\gamma' \geq \gamma$, and so uv is in the lower part of $A_{\gamma'}$ for all $\gamma' \geq \gamma$. It follows that there are no subsequent transfer steps between columns u and v.

Therefore $\{B_\gamma\colon \gamma \in Z\}$ is the sequence of arrays in an $(n-s,s)$-token game. Let e be an edge in G' with $\mathrm{drop}(G,S,e) = g(n,s)$, and let β be the index of the cell in A' containing e. By Lemma 3.2, we have that $g(n,s) = \mathrm{drop}(G,S,e) \le j-i$, where $[(i,u),(j,u)]$ is the critical interval of A_β. Since $A_\beta(\ell,u)$ contains a hole for $i \le \ell < j$, it follows that $B_\beta(\ell,u)$ contains a grounded token for $1 \le \ell \le j-i$. Hence, it is possible to place at least $j-i$ tokens in some column in an $(n-s,s)$-token game and so $\hat{g}(n-s,s) \ge j-i$. □

It remains to analyze the (n,s)-token game. Our main tool is to show that in an (n,s)-token game in which the number of tokens in a particular column grows substantially, it is possible to find a subgame with half the number of transfer steps in which a column gains a substantial number of tokens. Eventually, we obtain a contradiction since the number of tokens in a column cannot grow by more than the total number of transfer steps.

Lemma 3.4. *Let $B_0,\ldots,B_k$ be the transcript of a token game with a total of m tokens and at most 2^ℓ transfer steps. Suppose that some column initially contains a tokens in B_0 and ends with b tokens in B_k. If a' and r are integers such that $m < (a'+1)r/2$, then some subinterval of $B_0,\ldots,B_k$ is the transcript of a token game with m tokens and at most $2^{\ell-1}$ transfer steps, in which some column initially has at most a' tokens and ends with at least b' tokens, where $b' = b-a-r+1$.*

Proof. Choose j so that both $B_0,\ldots,B_j$ and $B_j,\ldots,B_k$ are transcripts of token games with at most $2^{\ell-1}$ transfer steps.

Let u index a column which initially has a tokens in B_0 and ends with b tokens in B_k, and let R be the set of tokens which end in column u but were not always in column u. Clearly, $|R| \ge b-a$. Let $\{t_1,\ldots,t_r\}$ be the tokens in R which are in the r highest positions in B_k, and let $R_0 = \{t_1,\ldots,t_r\}$. Each token in R_0 has height at least b' in B_k. Moreover, since the height of a token is non-increasing throughout the game, it follows that each token in R_0 has height at least b' in every array.

Since each $t_i \in R_0$ is moved to u during the token game, we may choose columns $v_1,\ldots,v_r$ and indices $\ell_1,\ldots,\ell_r$ such that t_i is moved from v_i in B_{ℓ_i} to u in B_{ℓ_i+1}. Since a token game forbids more than one transfer between a pair of columns, $v_1,\ldots,v_r$ are distinct. Choose $I \in \{[0,j],[j,k]\}$ so that $|R_1| \ge r/2$, where $R_1 = \{t_i \in R_0\colon \{\ell_i,\ell_i+1\} \subseteq I\}$. Since t_i has height at least b' throughout the game, column v_i in B_{ℓ_i} has at least b' grounded tokens.

Note that it is not possible for each of the columns in $\{v_i\colon t_i \in R_1\}$ to begin the subgame $\{B_i\colon i \in I\}$ with more than a' tokens, since $(a'+1)|R_1| \ge (a'+1)(r/2) > m$. It follows that some column v_i has at most a' tokens in the first array of $\{B_i\colon i \in I\}$ but has at least b' tokens in B_{ℓ_i}. □

Iterating Lemma 3.4 gives the following.

Lemma 3.5. *In a token game with m tokens and at most 2^ℓ transfer steps, each column gains a net of at most $1 + 2\ell\lceil\sqrt{2m}\rceil$ tokens.*

Proof. If $\ell = 0$, then the lemma is clear. For larger ℓ, suppose that there is a token game $B_0,\ldots,B_k$ with at most 2^ℓ transfer steps in which some column begins with a tokens and ends with b tokens. We iterate Lemma 3.4 with $a' = r = \lceil\sqrt{2m}\rceil$ to obtain, for each $1 \le t \le \ell$, a subgame with m tokens and at most $2^{\ell-t}$ transfer steps in which some column begins with at most a' tokens and ends with at least $(b-a)-(2t-1)r$ tokens.

With $\ell = t$, we obtain a token game with at most 1 transfer step in which some column begins with at most a' tokens and ends with at least $(b-a) - (2\ell - 1)r$ tokens. We conclude $(b-a) - (2\ell-1)r - a' \le 1$, which implies $b - a \le 1 + 2\ell r$. □

Corollary 3.6. *Always* $\hat{g}(n,s) \le 4 \lg n(\sqrt{2ns} + 1) + s + 1 = O(s + \sqrt{ns} \lg n)$. *In particular, if* $n \ge \max\{2, s\}$, *then* $\hat{g}(n,s) \le 11\sqrt{ns} \lg n$.

Proof. Consider an (n,s)-token game in which some column starts with at most s tokens and ends with $\hat{g}(n,s)$ tokens. Let $b = \hat{g}(n,s)$. Note that an (n,s)-token game contains at most 2^ℓ transfer steps provided that $2^\ell \ge \binom{n}{2}$; it suffices to choose $\ell = \lfloor 2 \lg n \rfloor$. Let m be the number of tokens in our (n,s)-token game; clearly $m \le ns$. It now follows from Lemma 3.5 that $b - s \le 1 + 2\ell\lceil\sqrt{2m}\rceil \le 1 + 4\lg n(\sqrt{2ns} + 1)$. When $n \ge \max\{2, s\}$, algebra gives the simpler bound $\hat{g}(n,s) \le 11\sqrt{ns} \lg n$. □

Improvements to Corollary 3.6 directly translate to improved bounds on $f(G)$ via Lemma 2.5 and Lemma 3.3. Unfortunately, our next theorem shows that there is not much room to improve Corollary 3.6.

Theorem 3.7. *Always* $\hat{g}(n,s) \ge \max\{s, \sqrt{2ns} - 3s/2\}$. *Consequently,* $\hat{g}(n,s) \ge \Omega(s + \sqrt{ns})$.

Proof. Clearly, $\hat{g}(n,s) \ge s$. Let k be the largest integer such that $s\binom{k+1}{2} \le n$ and note $k \ge \lfloor\sqrt{2n/s} - 1/2\rfloor$. Using that $n \ge s(1 + 2 + \cdots + k)$, we let $M_1, \ldots, M_k$ be disjoint sets of columns such that $|M_j| = sj$ for each j. Let $u_{j,1}, \ldots, u_{j,sj}$ be the columns in M_j. For each j, we construct a triangular pattern of tokens in M_j so that for $1 \le i \le sj$, the column $u_{j,i}$ contains i grounded tokens. We assume that the initial positions of all tokens are sufficiently high so that they fall into place as needed.

For M_1, we initialize the board so that for $1 \le i \le s$, the column $u_{1,i}$ starts with i tokens. For $j \ge 2$, we assume that we have played the token game so that for $1 \le i \le s(j-1)$, the column $u_{j-1,i}$ in M_{j-1} contains i grounded tokens. We use the tokens in M_{j-1} to construct the desired pattern in M_j. We move the highest grounded token from each column in $u_{j-1,1}, \ldots, u_{j-1,s(j-1)}$ to $u_{j,sj}$ in order. Since M_{j-1} has $s(j-1)$ columns, this puts $s(j-1)$ tokens in $u_{j,sj}$ and leaves a smaller triangular pattern in M_{j-1} where $u_{j-1,i}$ contains $i-1$ grounded tokens. Next, we move the highest grounded token from each column in $u_{j-1,2}, \ldots, u_{j-1,s(j-1)}$ to $u_{j,sj-1}$ in order; this places $s(j-1) - 1$ tokens in $u_{j,sj-1}$. Iterating this play, we move all tokens in M_{j-1} to M_j. We complete the triangular pattern by allowing $\min\{i, s\}$ tokens whose initial positions were high in column $u_{j,i}$ to fall into place.

Note that we require at most s tokens in each column initially. Moreover, since $M_1, \ldots, M_k$ are pairwise disjoint, no pair of columns is involved in more than one transfer step. After all steps, column $u_{k,sk}$ in M_k contains sk tokens, implying that $\hat{g}(n,s) \ge sk \ge s(\sqrt{2n/s} - 3/2)$. □

Although the bounds in Corollary 3.6 and Theorem 3.7 establish the order of growth of $\hat{g}(n,s)$ up to a logarithmic factor, it would still be interesting to obtain the exact order of growth. If $\hat{g}(n,s) = O(s + \sqrt{ns})$ as we suspect, then the log term in the lower bound in Corollary 4.2 can be removed. We do not know how sharp the inequality in Lemma 3.3 is; there may be room to make more substantial improvements to our upper bound on $g(n,s)$.

4 Dense Graphs

Theorem 4.1. *Let G be an n-vertex graph, let $s = n^{1/3}(11 \lg n)^{2/3}$, and suppose that n is sufficiently large so that $s \le n-2$. If G has average degree d and $d > 2$, then $f(G) > \frac{d}{4s}(1-\frac{2}{d})(1-\frac{1}{s})(1-\frac{4s^2}{d-2})$.*

Proof. Let H be a total ordering of G. Since H has $nd/2$ edges, it follows that some edge x_0x_1 has height at least $d/2$. With $s' = \lfloor s \rfloor$, we apply Lemma 2.5 to obtain a monotone path of length at least $s'\lfloor \frac{d/2-1}{\binom{s'+1}{2}+g(n,s')} \rfloor + 1$. Using Lemma 3.3, Corollary 3.6, and monotonicity of $\hat{g}(n,s)$, we have $g(n,s') \le \hat{g}(n-s',s') \le \hat{g}(n,s') \le 11\sqrt{ns'}\lg n$. We compute

$$\begin{aligned} s'\left\lfloor \frac{d/2-1}{\binom{s'+1}{2}+g(n,s')} \right\rfloor + 1 &> s'\left\lfloor \frac{d/2-1}{\binom{s'+1}{2}+11\sqrt{ns'}\lg n} \right\rfloor \\ &\ge (s-1)\left\lfloor \frac{d/2-1}{\binom{s+1}{2}+11\sqrt{ns}\lg n} \right\rfloor \\ &= (s-1)\left\lfloor \frac{d/2-1}{\binom{s+1}{2}+s^2} \right\rfloor \\ &\ge (s-1)\left(\frac{d/2-1}{2s^2}-1\right) \\ &= \frac{d}{4s}\left(1-\frac{2}{d}\right)\left(1-\frac{1}{s}\right)\left(1-\frac{4s^2}{d-2}\right) \end{aligned}$$

□

When the average degree d grows faster than s^2, Theorem 4.1 improves Rödl's result $f(G) \ge (1-o(1))\sqrt{d}$. In terms of n, an n-vertex graph must have average degree at least $Cn^{2/3}(\lg n)^{4/3}$ for some constant C in order for Theorem 4.1 to offer an improvement.

Corollary 4.2. $f(K_n) \ge (\frac{1}{20} - o(1))(\frac{n}{\lg n})^{2/3}$

Proof. By Theorem 4.1, we have $f(K_n) \ge \frac{n-1}{4n^{1/3}(11\lg n)^{2/3}}(1-o(1)) \ge \frac{n^{2/3}}{20(\lg n)^{2/3}}(1-o(1))$. □

We make no attempt to optimize the constant 1/20. Echoing remarks of Graham and Kleitman [6], we conjecture that our lower bound on $f(K_n)$ is not sharp and that the order of growth of $f(K_n)$ is closer to linear than to $(n/\log n)^{2/3}$.

References

[1] N. Alon. Problems and results in extremal combinatorics. I. *Discrete Math.*, 273(1-3):31–53, 2003. EuroComb'01 (Barcelona).

[2] A. P. Burger, E. J. Cockayne, and C. M. Mynhardt. Altitude of small complete and complete bipartite graphs. *Australas. J. Combin.*, 31:167–177, 2005.

[3] A. R. Calderbank, F. R. K. Chung, and D. G. Sturtevant. Increasing sequences with nonzero block sums and increasing paths in edge-ordered graphs. *Discrete Math.*, 50(1):15–28, 1984.

[4] V. Chvátal and J. Komlós. Some combinatorial theorems on monotonicity. *Canad. Math. Bull.*, 14:151–157, 1971.

[5] J. De Silva, T. Molla, F. Pfender, T. Retter, and M. Tait. Increasing paths in edge-ordered graphs: the hypercube and random graphs. *ArXiv e-prints*, Feb. 2015.

[6] R. L. Graham and D. J. Kleitman. Increasing paths in edge ordered graphs. *Period. Math. Hungar.*, 3:141–148, 1973. Collection of articles dedicated to the memory of Alfréd Rényi, II.

[7] J. Katrenič and G. Semanišin. Finding monotone paths in edge-ordered graphs. *Discrete Appl. Math.*, 158(15):1624–1632, 2010.

[8] M. Lavrov and P. Loh. Increasing hamiltonian paths in random edge orderings. *Random Structures & Algorithms*, 2015.

[9] C. M. Mynhardt, A. P. Burger, T. C. Clark, B. Falvai, and N. D. R. Henderson. Altitude of regular graphs with girth at least five. *Discrete Math.*, 294(3):241–257, 2005.

[10] Y. Roditty, B. Shoham, and R. Yuster. Monotone paths in edge-ordered sparse graphs. *Discrete Math.*, 226(1-3):411–417, 2001.

[11] V. Rödl. Master's thesis, Charles University, 1973.

[12] V. G. Vizing. On an estimate of the chromatic class of a p-graph. *Diskret. Analiz No.*, 3:25–30, 1964.

[13] P. Winkler. Puzzled: Solutions and sources. *Commun. ACM*, 51(9):103–103, 2008.

[14] R. Yuster. Large monotone paths in graphs with bounded degree. *Graphs Combin.*, 17(3):579–587, 2001.