

Pebbling and Optimal Pebbling in Graphs

David P. Bunde^{*} Erin W. Chambers[†] Daniel Cranston[‡]
Kevin Milans[§] Douglas B. West[¶]

Abstract

Given a distribution of pebbles on the vertices of a graph G , a *pebbling move* takes two pebbles from one vertex and puts one on a neighboring vertex. The *pebbling number* $\Pi(G)$ is the minimum k such that for every distribution of k pebbles and every vertex r , it is possible to move a pebble to r . The *optimal pebbling number* $\Pi_{OPT}(G)$ is the minimum k such that some distribution of k pebbles permits reaching each vertex.

We give short proofs of prior results on these parameters for paths, cycles, trees, and hypercubes, a new linear-time algorithm for computing $\Pi(G)$ on trees, and new results on $\Pi_{OPT}(G)$. If G is a connected n -vertex graph, then $\Pi_{OPT}(G) \leq \lceil 2n/3 \rceil$, with equality for paths and cycles. If \mathbf{G} is the family of n -vertex connected graphs with minimum degree k , then $2.4 \leq \max_{G \in \mathbf{G}} \Pi_{OPT}(G) \frac{k+1}{n} \leq 4$ when $k > 15$ and k is a multiple of 3. Finally, $\Pi_{OPT}(G) \leq 4^t n / ((k-1)^t + 4^t)$ when G is a connected n -vertex graph with minimum degree k and girth at least $2t+1$. For $t=2$, a more precise version of this last bound is $\Pi_{OPT}(G) \leq 16n/(k^2+17)$.

1 Introduction

Graph pebbling is a model for the transmission of consumable resources. Initially, pebbles are placed on the vertices of a graph G according to a *distribution* D , a function $D: V(G) \rightarrow$

^{*}Computer Science Department, University of Illinois, Urbana, IL 61801, bunde@uiuc.edu. Partially supported by NSF grant CCR 0093348.

[†]Computer Science Department, University of Illinois, Urbana, IL 61801, erinwolf@uiuc.edu. Supported by a National Science Foundation Graduate Research Fellowship.

[‡]Computer Science Department, University of Illinois, Urbana, IL 61801, dcransto@uiuc.edu.

[§]Computer Science Department, University of Illinois, Urbana, IL 61801, milans@uiuc.edu.

[¶]Mathematics Department, University of Illinois, Urbana, IL 61801, west@math.uiuc.edu. Work supported in part by the NSA under Award No. MDA904-03-1-0037.

$\mathbb{N} \cup \{0\}$. A *pebbling move* from a vertex v to a neighbor u takes away two pebbles at v and adds one pebble at u . Before the move, v must have at least two pebbles. A *pebbling sequence* is a sequence of pebbling moves.

Given a distribution and a “root” vertex r , the task is to put a pebble on r . A distribution D is *r -solvable* (and r is *reachable* under D) if r has a pebble after some (possibly empty) pebbling sequence starting from D . For a graph G , let $\Pi(G, r)$ be the least k such that every distribution of k pebbles on G is r -solvable. A distribution D is *solvable* if every vertex is reachable under D . The *pebbling number* of a graph G , denoted $\Pi(G)$, is the least k such that every distribution of k pebbles on G is solvable. The *optimal pebbling number* of G , denoted $\Pi_{OPT}(G)$, is the least k such that some distribution of k pebbles is solvable.

Graph pebbling originated in efforts of Lagarias and Saks to shorten a result in number theory. A survey by Hurlbert [7] describes this history and summarizes early results. Hurlbert introduced a useful generalization. A distribution D is *m -fold r -solvable* (and r is *m -reachable* under D) if r has at least m pebbles after some (possibly empty) pebbling sequence. A distribution D is *m -fold solvable* if every vertex is m -reachable under D . When $m = 2$, we say that an m -fold solvable distribution is *doubly solvable*.

Moews [9] developed several useful tools for computing pebbling numbers. (An unpublished longer version of the paper [9] appears on his webpage [11].) We call the first of these tools the *Weight Argument*, which we express here for m -fold solvability. Given a root r and distribution D , let $a_{i,r}$ be the total number of pebbles on vertices at distance i from r . A pebbling move cannot increase the sum $\sum_{i \geq 0} a_{i,r} 2^{-i}$. Therefore, m -fold r -solvability of D requires the *weight inequality* $\sum_{i \geq 0} a_{i,r} 2^{-i} \geq m$.

Our other main tool is that when each pebbling move is represented by a directed edge from the vertex losing pebbles to the vertex gaining a pebble, no directed cycle is needed. If r is reachable using moves containing a cycle, then also r is reachable using a proper subset of these moves. In particular, if a distribution is r -solvable, then r is reachable without moving a pebble in both directions along any edge.

To make this precise, say that a directed multigraph H is *orderable* under a distribution D if some linear ordering σ of $E(H)$ is a valid pebbling sequence starting from D . For such D and H , the *balance* of a vertex v is $d_H^-(v) + D(v) - 2d_H^+(v)$, where $d_H^-(v)$ and $d_H^+(v)$ are the indegree and outdegree of v under H . When H is orderable under D (by σ), each vertex has nonnegative balance, since the balance is the number of pebbles at v after applying σ . The *No-Cycle Lemma* states that if H is orderable under D , then it has an acyclic subgraph H' such that H' is orderable under D and gives balance to each vertex at least as large as

does H . The lemma was proved in [3] and earlier in [9] and has a short proof in [8].

The pebbling number is known exactly for some special graphs. Moews [9] observed that a distribution on a path rooted at its end is solvable if and only if the weight inequality holds; thus $\Pi(P_n) = 2^{n-1}$ for the n -vertex path P_n , since each pebble contributes weight at least $2^{-(n-1)}$. The n -vertex cycle C_n is more complicated; Pachter et al [12] proved that $\Pi(C_{2k}) = 2^k$ and $\Pi(C_{2k+1}) = 2 \lfloor 2^{k+1}/3 \rfloor + 1$. For the k -dimensional hypercube Q_k , Chung [4] proved that $\Pi(Q_k) = 2^k$. For a rooted tree, Moews [9] showed how to calculate the pebbling number from decompositions into paths. For general graphs, Milans and Clark [8] showed that recognizing $\Pi(G) \leq k$ is a Π_2^P -complete problem, meaning that it is complete for the class of languages computable in polynomial time by coNP machines equipped with an oracle for an NP-complete language.

Study of the optimal pebbling number began with the result of Pachter et al [12] that $\Pi_{OPT}(P_n) = \lceil 2n/3 \rceil$. Moews [10] proved that $(4/3)^k \leq \Pi_{OPT}(Q_k) \leq (4/3)^{k+O(\log k)}$ and proved a related result for Π_{OPT} on cartesian product graphs. Milans and Clark [8] proved that computing Π_{OPT} is NP-hard on arbitrary graphs.

In this paper, we present several new results (mostly on optimal pebbling) and simpler proofs for some previously-known results. The proofs are simple due to lemmas that restrict the form of pebble distributions that need to be considered. These lemmas and another lower-bound technique for optimal pebbling comprise our main contribution. Our undirected graphs are simple and connected. We use $V(G)$ and $E(G)$ to denote the vertex set and edge set of a graph G , with sizes $n(G)$ and $e(G)$.

For the pebbling number, we give another proof of the result of Moews [9] on calculating $\Pi(T, r)$ from a particular decomposition of a tree T into paths. We extend his result to give a linear-time algorithm for computing $\Pi(T)$. Also, we give short proofs of the results of Pachter et al [12] that $\Pi(C_{2k}) = 2^k$ and $\Pi(C_{2k+1}) = 2 \lfloor 2^{k+1}/3 \rfloor + 1$.

Our approach (especially for paths and cycles), relies on a precise version of the following intuition. For distributions with k pebbles, the hardest ones to make solvable are concentrated on one or two vertices, while the easiest ones are spread over many vertices. Thus to determine $\Pi(G)$ we consider concentrated distributions, while to determine $\Pi_{OPT}(G)$ we consider “smooth” distributions.

For optimal pebbling, the *Smoothing Lemma* is that for each graph a solvable distribution of minimum size exists with at most two pebbles on each vertex of degree at most 2. This leads to a simpler proof of the result of Pachter et al [12] that $\Pi_{OPT}(P_n) = \lceil 2n/3 \rceil$ and a proof that $\Pi_{OPT}(C_n) = \lceil 2n/3 \rceil$. (We recently learned that Friedman and Wyels [6] have

obtained another short derivation of $\Pi_{OPT}(P_n)$ different from ours, and like us they adapted those ideas to compute $\Pi_{OPT}(C_n)$.)

We also show that $\Pi_{OPT}(T) \leq \lceil 2n/3 \rceil$ for every n -vertex tree T , which immediately yields $\Pi_{OPT}(G) \leq \lceil 2n/3 \rceil$ for every connected n -vertex graph G , and we give a short proof of the result of Moews [10] that $\Pi_{OPT}(Q_k) \geq (4/3)^k$.

Let G be a connected n -vertex graph with minimum vertex degree k . Czygrinow [5] observed that $\Pi_{OPT}(G) \leq 4\frac{n}{k+1}$. We construct families of such graphs with $\Pi_{OPT}(G) \geq 2\frac{n}{k+1}$ for all k and with $\Pi_{OPT}(G) \geq (2.4 - \frac{24}{15k+5})\frac{n}{k+1}$ when k is divisible by 3. These results use another lower-bound technique, the simplest version of which is that if G is obtained from H by collapsing sets of vertices into single vertices, then $\Pi_{OPT}(H) \geq \Pi_{OPT}(G)$.

We obtain tighter bounds when we further restrict G to have girth (minimum cycle length) at least $2t + 1$. Suppose that $k \geq 3$ and $t \geq 2$, but exclude the case $(k, t) = (3, 2)$. Letting $c_k(t) = 1 + k \sum_{i=1}^t (k-1)^{i-1}$ and $c'(t) = (4^t - 2^{t+1})\frac{t}{t-1}$, we prove that $\Pi_{OPT}(G) \leq 4^t n / (c_k(t) + c'(t))$. When G has girth at least 5, this yields $\Pi_{OPT}(G) \leq \frac{16n}{k^2+17}$. Among graphs with girth 4, we show that $\Pi_{OPT}(C_m \square K_2) = \Pi_{OPT}(P_m \square K_2) = m$ when $m \geq 3$ (except $m+1$ when $m=5$), where \square denotes cartesian product (see Section 6). The same bound holds also for the graph consisting of a $2m$ -cycle with chords added joining opposite vertices.

Our results on pebbling number of trees and pebbling number of cycles appear in Sections 2 and 3, respectively. The final three sections discuss optimal pebbling number and present the new results. In addition to the results mentioned above, we pose the question of whether every connected n -vertex graph with minimum degree at least 3 has optimal pebbling number at most $\lceil n/2 \rceil$.

2 Pebbling Number of Trees

Moews [9] showed how to compute the pebbling number of a tree from a decomposition into paths. In this section, we prove this more simply and show how to find an optimal decomposition in linear time.

A partition of the edge set of a tree is a *path partition* if each set in the partition is a (directed) path when all edges are directed toward a root r . The *length list* of a path partition is the list of path lengths in nonincreasing order. A path partition *majorizes* another if its length list is larger than the other's in the first position where they differ. Majorization is a linear (lexicographic) order on length lists, but distinct path partitions may have length lists that are the same. A path partition with root r is *r-optimal* if it is not majorized by any

other path partition with root r . It is *optimal* if it is not majorized by any path partition with any root. We use *leaf* to describe a vertex of degree 1 in any graph.

Moews [9] showed how to determine $\Pi(T, r)$ from an optimal path partition of a tree T rooted at a vertex r . Our proof is shorter and simpler.

Theorem 1 (Moews [9]). *If the length list of an r -optimal path partition of tree T with root r is l_1, \dots, l_m , then*

$$\Pi(T, r) = \left(\sum_{i=1}^m 2^{l_i} \right) - m + 1.$$

Proof. We have observed r -solvability never requires moving pebbles in both directions along an edge. Thus in a tree we may direct all edges toward the root and move pebbles only in that direction. Let \mathcal{L} be an optimal path partition of a tree T rooted at r , and let (l_1, \dots, l_m) be the length list of \mathcal{L} .

Lower Bound. We construct a non- r -solvable distribution with $\sum_{i=1}^m (2^{l_i} - 1)$ pebbles. If some path in \mathcal{L} starts at a nonleaf vertex, then another path ends there, and they combine to produce a path partition majorizing \mathcal{L} . Hence in \mathcal{L} each path begins at a leaf. For each path of length l_i in \mathcal{L} , we put $2^{l_i} - 1$ pebbles on the starting leaf. Now no pebble can be the first pebble to reach the end of the path in \mathcal{L} on which it starts. Hence pebbles never reach the end or move off their starting path in \mathcal{L} . In particular, no pebble can reach r .

Upper Bound. We show that every distribution with more than $(\sum_{i=1}^m 2^{l_i}) - m$ pebbles is r -solvable, using a weight function based on \mathcal{L} . Let P_i be the path in \mathcal{L} corresponding to length l_i . Given a distribution D , let $a_{i,j}$ be the number of pebbles on P_i at distance j from the end. Let $w_i(D) = 2^{l_i} \sum_{j=1}^{l_i} a_{i,j} 2^{-j}$, and let $w(D) = \sum_{i=1}^m w_i(D)$.

The function $w_i(D)$ differs from the standard weight function on a path in two ways: we multiply by an extra factor of 2^{l_i} , and we sum over $j \geq 1$ rather than $j \geq 0$. We sum over $j \geq 1$ because the end of P_i is inside a longer path (or is r); we avoid counting pebbles twice. The factor of 2^{l_i} ensures that moves toward r do not decrease the total weight.

A path P_i in \mathcal{L} is *full* under distribution D if $w_i(D) \geq 2^{l_i}$. If P_i is full, then we can move a pebble along P_i to its end, where it will be on r or contribute to the weight of a path that extends closer to r . Each move within a path does not change the total weight. When a pebble moves from P_i to $P_{i'}$, the weight decreases by $2 \cdot 2^{l_i - 1}$ and increases by $2^{l_{i'} - j}$, where j is the distance from the new location to the end of $P_{i'}$. If $l_i > l_{i'} - j$, then P_i can replace the beginning of $P_{i'}$ to produce a path partition majorizing \mathcal{L} ; the optimality of \mathcal{L} prevents this. Hence $l_i \leq l_{i'} - j$, and moving a pebble from P_i to $P_{i'}$ does not decrease the weight.

Given an optimal path partition with lengths l_1, \dots, l_m , let D be a distribution under which r is not reachable. If D has more than $\sum_{i=1}^m (2^{l_i} - 1)$ pebbles, then by the pigeonhole principle some path is full, since each pebble not on r contributes at least 1 to the weight of the tree. We have shown that no move toward r decreases the total weight, except when a pebble is moved onto r and no longer contributes. Every pebbling sequence terminates, since each move reduces the total number of pebbles. Since the weight never decreases, the sequence can only terminate by moving a pebble onto r . \square

In his survey [7], Hurlbert attributes the corollary to Moews.

Corollary 2. *If the length list of an optimal path partition of tree T is l_1, \dots, l_m then*

$$\Pi(T) = \sum_{i=1}^m 2^{l_i} - m + 1.$$

Proof. Since exponentiation is a convex function, the formula in Theorem 1 is maximized by an r -optimal path partition. Also $\Pi(T) = \max_{r \in V(T)} \Pi(T, r)$. Hence the claim follows from Theorem 1. \square

The difficulty in applying Corollary 2 is in finding an optimal path partition. Given a root, a natural idea is to select a longest path greedily and iterate. Although this works, it disconnects the tree, leaving awkward bookkeeping details. The inductive proof is simpler if we peel away shorter paths first. A *peripheral vertex* in a tree is an endpoint of a longest path. A *branch vertex* in a tree is a vertex of degree at least 3. An x, y -*path* in a graph is a path with endpoints x and y .

Theorem 3. *There is a linear-time algorithm to compute the pebbling number of trees. In particular, if r is an endpoint of a longest path in T , then $\Pi(T, r) = \Pi(T)$, and any longest path to r can be chosen as a path in an r -optimal path partition.*

Proof. In a tree, the vertices at greatest distance from a vertex x are endpoints of a longest path. Hence a single breadth-first search from an arbitrary vertex finds a peripheral vertex r . Another breadth-first search from r finds a longest path R , ending at another vertex r' .

With R chosen, another breadth-first search computes distances from R . We find an r -optimal path partition using these distances. The partition will have R as a path, and it will be both r -optimal and r' -optimal. We view all edges off R as directed toward R .

Suppose that R is not all of T . Iteratively, we select a leaf x closest to R among the leaves that remain in the tree. Let y be the closest branch vertex to x in T ; vertex y is well-defined. Since R is a longest path, y cannot be r or r' . Let P be the x, y -path in T . Put P into the path partition and delete P from the tree, leaving only the endpoint y . When the remaining tree is just R , it becomes the last path in the partition. (We can pause the computation of distances from R each time a leaf is found and extract P then.)

We prove, by induction on the number of vertices outside R , that the path P deleted at each step lies in an r -optimal path partition of the tree remaining at that step. By the majorization criterion, the path P' containing x in an r -optimal path partition \mathcal{L} contains all of P . If P' continues past y , then some path Q in \mathcal{L} ends at y . We have observed that Q starts at a leaf, so Q is at least as long as P , by the choice of x .

Let Q' be the union of Q and the part of P' after y . Let \mathcal{L}' be the partition obtained from \mathcal{L} by replacing P' and Q with P and Q' . Now P and Q' are shortest and longest, respectively, among $\{P, P', Q, Q'\}$. If Q is longer than P , then \mathcal{L}' majorizes \mathcal{L} . Otherwise, \mathcal{L}' and \mathcal{L} have the same length list; hence \mathcal{L} is an r -optimal path partition containing P .

Thus P occurs in an r -optimal path partition \mathcal{L} . The remainder of \mathcal{L} is an r -optimal path partition of the remaining tree T' . Distances from R are the same in T' as in T . By the induction hypothesis, the remainder of the algorithm produces an r -optimal path partition of T' that contains R . It combines with P to yield the desired path partition of T .

The partition we have produced is also r' -optimal, since the computation is the same when viewed from r' (distances from R are the same).

Since we have found an r -optimal path partition containing a longest path, the length list of a globally optimal path partition must include the longest path length. Hence $\Pi(T)$ equals $\Pi(T, r)$ for some peripheral vertex r .

We show next that the procedure produces the same length list from each peripheral vertex. When r and r' are the endpoints of a longest path R , we showed that r -optimal and r' -optimal path partitions have the same length list. When R' is another longest path from r' , the algorithm would again produce an r' -optimal path partition. Since the lexicographic order is linear, all r' -optimal path partitions have the same length list.

Since every longest path in a tree contains the center of the tree, if the path joining two peripheral vertices is not a longest path, then each is an endpoint of a longest path to one other peripheral vertex. Hence one can move from one peripheral vertex to any other by at most two instances of “move to the opposite end of a longest path”. Therefore, all peripheral vertices have the same optimal length list.

Because an optimal path partition must contain a longest path and hence must be an r -optimal path partition for some peripheral vertex r , we conclude that $\Pi(T) = \Pi(T, r)$ for each peripheral vertex r . \square

3 Pebbling Number of Cycles

Proving an upper bound on the pebbling number requires showing that each of a large number of distributions is solvable. The following lemma restricts the distributions that need to be considered. A *thread* in a graph G is a path whose vertices have degree 2 in G .

Lemma 4 (Squishing Lemma). *For a vertex r in a graph G , there is a non- r -solvable distribution of $\Pi(G, r) - 1$ pebbles on G such that on each thread not containing r , all pebbles occur on just one vertex or on two adjacent vertices.*

Proof. Let P be a thread in G . If a distribution has pebbles on only one vertex of P or on only two adjacent vertices of P , then we say that P is *squished*.

Let D be a distribution of $\Pi(G, r) - 1$ pebbles that is not r -solvable. We transform D into a distribution of the same size such that every thread not containing r is squished. A *squishing move* removes 1 pebble from each of two vertices on a thread and puts 2 pebbles on some vertex between them on the thread. If some path P is not squished, then we can perform a squishing move on P . Each squishing move reduces the value of $\sum_p 2^{-b(p)}$, where the sum is over the set of pebbles on P and $b(p)$ is the distance of pebble p from a fixed end of P . Thus a sequence of squishing moves must end by squishing P .

Let D' be the result of a squishing move applied to D on a thread of P not containing r ; pebbles from y and z are moved to x between them. We show that if D' is r -solvable, then D is r -solvable. Let σ be a pebbling sequence from D' that reaches r . If σ never moves pebbles off x , then σ also reaches r from D . Hence we may assume that σ includes a move from x to a neighbor x' , which we may assume is toward y along P .

By the No-Cycle Lemma, we may assume that σ makes no move from x' to x . The two pebbles used to move from x to x' thus produce no more benefit than the one pebble that started on y in D ; under D starts farther than x' in the only direction it can go. Also it cannot hurt to have the extra pebble on z . Thus D also is r -solvable. \square

The Squishing Lemma provides a short proof for the pebbling number of C_n .

Theorem 5 (Pachter et al [12]). *The pebbling number of the cycle satisfies $\Pi(C_{2k}) = 2^k$ and $\Pi(C_{2k+1}) = 2 \lfloor 2^{k+1}/3 \rfloor + 1$.*

Proof. *Lower Bound.* Given a root r in C_{2k} , a distribution with $2^k - 1$ pebbles on the one vertex at distance k from r is not r -solvable. We show that in C_{2k+1} , a distribution with $\lfloor 2^{k+1}/3 \rfloor$ pebbles on each of the two vertices at distance k from r is not r -solvable. One pile alone cannot move distance k to reach r . If we combine them first, moving half of one pile to the other, then the resulting pile has at most $\frac{2^{k+1}-1}{3} + \frac{1}{2} \frac{2^{k+1}-1}{3}$ pebbles, since 2^{k+1} is not divisible by 3. The sum is less than 2^k , so again the pile cannot reach r .

Upper Bound. A distribution having 2^k pebbles on some path of length k ending r is r -solvable, since $\Pi(P_{k+1}) = 2^k$. This suffices for most cases, since the Squishing Lemma allows us to restrict attention to distributions covering only one or two adjacent vertices. In C_{2k} , every two adjacent vertices lie together in a path of length k ending at r . This also holds for all cases in C_{2k+1} except when the two adjacent vertices are the two vertices s and s' at distance k from r .

In this case, with all the pebbles on $\{s, s'\}$, we move as many as possible from the vertex with fewer pebbles to the vertex with more pebbles. With m pebbles total and l in the smaller pile, the new pile has size at least $m - l + \lfloor l/2 \rfloor$. Since $l \leq \lfloor m/2 \rfloor$ and $m \geq 2 \lfloor 2^{l+1}/3 \rfloor + 1$, we obtain a pile of size at least 2^k at distance k from r , which suffices. \square

4 Optimal Pebbling Number

For optimal pebbling numbers, upper bounds are generally easier than lower bounds. For an upper bound, we give a distribution and show that it is solvable. For a lower bound, we must show that every distribution up to a certain size is not solvable.

The Smoothing Lemma plays the role for optimal pebbling that the Squishing Lemma plays for ordinary pebbling. The purpose again is to restrict the form of distributions we study to determine the value of the parameter. Instead of squishing pebbles together on a thread, we spread them out.

When D is a distribution on a graph with a vertex v of degree 2, and v has at least three pebbles in D , a *smoothing move* from v changes D by removing two pebbles from v and adding one pebble at each neighbor of v . The case $m = 2$ below will be used in Section 5.

Lemma 6. *Let D be a distribution on a graph G with distinct vertices u and v , where v has degree 2. If $D(v) \geq 3$, and u is m -reachable under D , then u is m -reachable under the distribution D' obtained by making a smoothing move from v .*

Proof. For any pebbling sequence σ starting from D , we form a sequence σ' from D' . If σ never makes a move from v , then we may set $\sigma' = \sigma$, since at each step there are at least as many pebbles at each vertex other than v when starting with D' .

If σ makes a move from v , then let σ' be the same as σ except that σ' skips the first such move. Having made that move, σ on D produces the same configuration as σ' on D' , except that σ' on D' has an extra free pebble on one neighbor of v . We complete σ' using the rest of σ and have the same number of pebbles at each vertex as under σ from D , plus an extra pebble on one neighbor of v . (Since σ' mimics σ , we never use that extra pebble.) \square

A distribution D is *smooth* if it has at most two pebbles on every vertex of degree 2 (so no smoothing move is possible). A vertex D is *unoccupied* under D if $D(v) = 0$.

Lemma 7 (Smoothing Lemma). *If G is connected and $n(G) \geq 3$, then G has a smooth minimal solvable distribution with all leaves unoccupied.*

Proof. A minimal solvable distribution has $\Pi_{OPT}(G)$ pebbles, and always $\Pi_{OPT}(G) \leq n(G)$. We first transform an arbitrary solvable distribution D with $|D| \leq n(G)$ into a smooth solvable distribution of the same size; later we also eliminate pebbles from leaves.

By Lemma 6, a smoothing move from v preserves the reachability of vertices other than v . Since a smoothing move from v leaves a pebble at v , also v remains reachable. Therefore, smoothing moves preserve solvability. To complete the proof of the first claim, it suffices to show that a smooth distribution will result from applying smoothing moves to any distribution with at most $n(G)$ pebbles.

Suppose first that G is not a cycle. Starting from any distribution on G , we show that only finitely many smoothing moves can be made. Every vertex v of degree 2 lies in a unique maximal thread. Let P be the unique path through v whose internal vertices have degree 2 and whose endpoints do not. When P has length m and v has distance k from one end of P , we count each pebble on v with weight $k(m - k)$; it does not matter which end the distance is measured from. Pebbles on a vertex with degree other than 2 count with weight 0.

Let v be a vertex at distance k from the end of a thread of length m (here the ends have degree other than 2). A smoothing move from v replaces weight $2k(m - k)$ at v with weight

$(k-1)(m-k+1) + (k+1)(m-k-1)$ at its neighbors. The total weight declines by 2. It must remain nonnegative, so we reach a distribution with no smoothing move available.

When G is a cycle, we use induction on the number of unoccupied vertices. Since $|D| \leq n(G)$, when all vertices are occupied there is one pebble on each vertex and D is smooth. If v has no pebbles and D is not smooth, then we view v as both endpoints of a thread around the cycle. Using the same weight argument as above, each smoothing move reduces the total weight by 2. Thus eventually the distribution becomes smooth or a pebble moves to v . Since smoothing never uncovers a vertex, moving a pebble to v reduces the number of unoccupied vertices. Thus the continuation of the smoothing process produces a smooth distribution.

We have obtained a smooth minimal solvable distribution D ; now we consider leaves. Let v be a leaf, and let u be its neighbor. Suppose that $D(u) = j$ and $D(v) = k \geq 1$.

Case 1: $j + k \geq 3$. Modify D by deleting the pebbles on v and adding $k - 1$ pebbles to u instead. The resulting D' is still solvable, since $D'(u) \geq 2$ makes v reachable, and D' starts with at least as many pebbles on u as v could send there to help pebble other vertices. However, $|D'| < |D|$, which contradicts the minimality of D .

Case 2: $j + k = 2$. Modify D by putting both pebbles on u . Still D' is smooth if u has degree 2. The two pebbles can be used to cover v , and they provide as much help for other vertices as before.

Case 3: $(j, k) = (0, 1)$. Move the one pebble to u ; again D' is smooth. Because D is u -solvable and cannot use the pebble on v to reach u , we can now move another pebble to u and use the two of them to reach v . □

The Smoothing Lemma yields a short proof of the result of Pachter et al [12] that $\Pi_{OPT}(P_n) = \lceil 2n/3 \rceil$, and it yields the same value also for cycles. Another short proof was given by Friedman and Wyels [6]. We separate an observation useful in Section 6.

Lemma 8. *Let v be an unoccupied vertex in a smooth distribution D on a path with at most two pebbles on each endpoint. If v is an endpoint, then v is not 2-reachable under D . If v is an internal vertex, then no pebbling sequence can move a pebble out of v without using an edge in both directions.*

Proof. The first claim follows immediately from the case $m = 2$ of the Weight Argument, since each vertex has at most two pebbles. For the second claim, moving a pebble out of v without first moving a pebble in from each neighbor would require contradicting the first

claim on a smaller path. □

Theorem 9. $\Pi_{OPT}(C_n) = \Pi_{OPT}(P_n) = \lceil 2n/3 \rceil$.

Proof. Let G be C_n or P_n .

Upper Bound. Partition G into $\lfloor n/3 \rfloor$ copies of P_3 and possibly one or two leftover vertices. Put two pebbles on the central vertex of each P_3 and one pebble on each of the leftover vertices (if any exist). The distribution is solvable and has size $\lceil 2n/3 \rceil$.

Lower Bound. By Lemma 7, it suffices to consider a smooth solvable distribution D with no pebbles on leaves. We use induction on n , checking $n \leq 5$ exhaustively.

By the No-Cycle Lemma, we may assume that the directed edges representing moves in a pebbling sequence to reach a target vertex form edge-disjoint paths, and no edge is used in both directions. Since D is smooth, Lemma 8 implies that each such path has no unoccupied internal vertex.

Since $\Pi_{OPT}(G) \leq \lceil 2n/3 \rceil$ and $n \geq 6$, at least two vertices of G are unoccupied. We may choose three unoccupied vertices, since otherwise $n = 6$, no vertex has two pebbles, and D is not solvable. With three unoccupied vertices, we can choose an unoccupied internal vertex in P_n or nonadjacent unoccupied vertices in C_n ; let S be this chosen set.

Since pebbles cannot be sent across an unoccupied vertex, S splits G into two paths, each of which cannot contribute pebbles to help pebble a vertex on the other path. Since the distribution is solvable, each vertex of S can be pebbled; we treat the vertex as being part of the path that pebbles it, choosing one such path if both can pebble it.

We now have paths of order l and $n - l$ with $1 \leq l \leq n - 1$, and D breaks into solvable distributions for these two paths. By the induction hypothesis, the number of pebbles in D is at least $\lceil 2l/3 \rceil + \lceil 2(n - l)/3 \rceil$, which is at least $\lceil 2n/3 \rceil$. □

Next we show that the path is a hardest tree for optimal pebbling number. It is far from unique; there are many trees whose optimal pebbling number is $\lceil 2n/3 \rceil$. We write $d(v)$ for the degree of a vertex v , and $N(v)$ for the set of vertices adjacent to v .

Theorem 10. *If T is an n -vertex tree, then $\Pi_{OPT}(T) \leq \lceil 2n/3 \rceil$.*

Proof. We use induction on n . The claim holds for $n \leq 3$, since all such trees are paths. In the induction step ($n > 3$), we delete three or more vertices at or near the end of a longest path in T to obtain a subtree T' . It suffices to show that we can add two pebbles to a

minimal solvable distribution D' on T' to form a solvable distribution D on T . When we add pebbles to D' , all vertices in T' remain reachable, so the problem reduces to showing that the new vertices can be reached.

Let P be a longest path in T . Let z be an endpoint of P , adjacent to y , and let x be the other neighbor of y on P . We consider four cases.

Case 1: $d(y) > 2$. Since P is a longest path, all neighbors of y other than x are leaves. Let $T' = T - y - (N(y) - \{x\})$. Form D from D' by adding two pebbles on y ; these make leaf neighbors of y reachable.

Case 2: $d(x) = d(y) = 2$. Let $T' = T - \{x, y, z\}$. Form D from D' by adding two pebbles on y ; these make x and z reachable.

Case 3: $d(y) = 2$ and x has a leaf neighbor u . Let $T' = T - \{u, y, z\}$. Form D from D' by adding two pebbles on y . Now y and z are reachable. We can also reach u by moving a pebble to x using the distribution D' on T' and then moving a second pebble to x from y .

Case 4: $d(y) = 2$ and x has no leaf neighbors. Let u be a neighbor of x outside P . Since P is a longest path, every neighbor of u other than x is a leaf. Let v be a leaf neighbor of u , and let $T' = T - \{v, y, z\}$. If x is 2-reachable under D' , then we form D by adding two pebbles on x , making $\{v, y, z\}$ all reachable under D . If u is 2-reachable under D' , then v is reachable, so we form D by adding two pebbles on y . If neither x nor u is 2-reachable under D' , then no pebbling sequence starting with D' uses the edge xu in either direction. Hence from D' we can reach x and u simultaneously. Now we form D by adding two pebbles on y , making $\{v, y, z\}$ all reachable after moving pebbles to both x and u using D' . \square

Corollary 11. *If G is a connected n -vertex graph, then $\Pi_{OPT}(G) \leq \lceil 2n/3 \rceil$, which is sharp.*

Proof. Adding an edge to a graph cannot increase its optimal pebbling number. Since G is connected, it has a spanning tree T . Applying Theorem 10 to T gives the bound, which is achieved by P_n . \square

Finally, we give a short proof that $\Pi_{OPT}(Q_k) \geq (4/3)^k$. The proof by Moews [10] used a continuous relaxation of pebbling, but the standard weight function and expectation suffice.

Theorem 12 (Moews [10]). $\Pi_{OPT}(Q_k) \geq (4/3)^k$, where Q_k is the k -dimensional hypercube.

Proof. Let D be a solvable distribution on Q_k ; we show that $|D| \geq (4/3)^k$. Since D is solvable, the standard weight inequality $\sum_{i \geq 0} a_{i,r} 2^{-i} \geq 1$ holds for each vertex r , where $a_{i,r}$ is the number of pebbles at distance i from r in D .

Select a vertex r in Q_k uniformly at random. Since the weight inequality holds for each r , linearity of expectation yields $\sum_{i \geq 0} 2^{-i} \mathbf{E}[a_{i,r}] \geq 1$. For a fixed pebble on a vertex u , the probability that r has distance i from u is $\binom{k}{i} 2^{-k}$, since Q_k has 2^k vertices and $\binom{k}{i}$ of them have distance i from r . By linearity of expectation, $\mathbf{E}[a_{i,r}] = |D| \binom{k}{i} 2^{-k}$. Substituting and simplifying now yields

$$|D| \sum_{i \geq 0} \binom{k}{i} 2^{-i} \geq 2^k.$$

Applying the Binomial Theorem yields $|D|(1 + \frac{1}{2})^k \geq 2^k$, and hence $|D| \geq (4/3)^k$. \square

5 Bounds in Terms of Minimum Degree

We have proved that $\Pi_{OPT}(G) \leq \lceil 2n/3 \rceil$ for every connected n -vertex graph G , with equality for paths and cycles. One would expect that tighter upper bounds hold for denser graphs. How large can $\Pi_{OPT}(G)$ be when we require minimum degree k ?

A *dominating set* in a graph G is a set $S \subseteq V(G)$ such that every vertex not in S has a neighbor in S . The *domination number* $\gamma(G)$ is the minimum size of a dominating set. Placing two pebbles at each vertex of a dominating set yields $\Pi_{OPT}(G) \leq 2\gamma(G)$. Thus upper bounds on $\gamma(G)$ yield upper bounds on $\Pi_{OPT}(G)$.

For graphs with minimum degree at least k , Arnaudov [2] and Payan [13] proved that $\gamma(G) \leq n \frac{1 + \ln(k+1)}{k+1}$; a short probabilistic argument appears in Alon [1]. In a k -regular n -vertex graph, dominating sets have size at least $\frac{n}{k+1}$, and Alon [1] showed that the domination number may be as large as $(1 + o(1))n \frac{1 + \ln(k+1)}{k}$. Hence we cannot improve the bound using domination number alone.

Czygrinow [5] communicated to us an easy argument for a better upper bound when $k \geq 3$; we begin by presenting this. A *distance-2 dominating set* in a graph G is a set $S \subseteq V(G)$ such that every vertex of G is within distance at most 2 from S (similarly, one can define distance- d dominating sets). The case $d = 1$ of the following proposition is folklore in some circles but seems to be unknown in the subject of graph domination. We will use the general result in Section 6.

Proposition 13. *If c is the minimum size of a distance- d neighborhood in G , then G has a distance- $2d$ dominating set of size at most $n(G)/c$.*

Proof. We build such a set S . Initially, put one vertex in S . As we proceed, let T consist of all vertices within distance d of S . If T is not a distance- d dominating set, then let v be a vertex that is not within distance d of T . Add v to S ; this adds the distance- d neighborhood of v to T , none of which was in T before. Thus T grows by at least c vertices for each vertex added to S . We therefore add at most n/c vertices to S by the time T becomes a distance- d dominating set, at which point S is a distance- $2d$ dominating set. \square

Corollary 14 (Czygrinow). *If G is a graph with minimum degree k , then $\Pi_{OPT}(G) \leq \frac{4n(G)}{k+1}$.*

Proof. Distance-1 neighborhoods have size at least $k+1$, so Proposition 13 yields a distance-2 dominating set S of size at most $n(G)/(k+1)$. Put four pebbles at each vertex of S . \square

Corollary 14 improves the upper bound of $\lceil 2n/3 \rceil$ from Corollary 11 when $k \geq 6$. A simple construction shows that this easy upper bound is within a factor of 2 of being sharp; we present n -vertex graphs with minimum degree k whose solvable distributions have at least $\frac{2n}{k+1}$ pebbles. Subsequently we present a better construction with optimal pebbling number approximately $\frac{2.4n}{k+1}$.

We begin by introducing another technique for proving lower bounds. Given a graph G , the operation of *collapsing* a vertex set $S \subseteq V(G)$ produces a new graph H in which S is replaced with a single vertex whose neighbors are the neighbors of S in G that were outside S . The subgraph induced by $V(G) - S$ remains unchanged. We use the term “collapsing” rather than “contracting” because the subgraph of G induced by S need not be connected.

Lemma 15 (Collapsing Lemma). *If H is obtained from G by collapsing vertex sets, then $\Pi_{OPT}(G) \geq \Pi_{OPT}(H)$.*

Proof. Let D be a solvable distribution on G . Form distribution D' on H as follows: for each collapsed set S , put all the pebbles that were on S in D onto the single vertex representing S in H . Treat uncollapsed vertices as collapsed sets of size 1.

To show that D' is solvable, for $u \in V(H)$ choose a vertex $v \in V(G)$ in the set that collapses to u . Let σ be a pebbling sequence from D that reaches v . The sequence σ “collapses” in an obvious way to a sequence σ' from D' that reaches u . More precisely, the distribution C resulting from a pebbling move on D collapses to a distribution C' on H that is obtained from D' by discarding one pebble (if the move on D was within a collapsed set) or by making one pebbling move from D' . \square

Proposition 16. *For $n > k \geq 2$, there is an n -vertex graph G with minimum degree k such that $\Pi_{OPT}(G) > \frac{2n}{k+1} - 2$, improving to $\Pi_{OPT}(G) \geq \frac{2n}{k+1}$ when n is a multiple of $k+1$.*

Proof. When $n = k + 1$, the complete graph K_n has this behavior.

When n is a larger multiple of $k + 1$, let J be the graph obtained from K_{k+1} by deleting one edge; the *internal* vertices of J are its vertices of degree k . Let G be the k -regular “ring of cliques” with $r(k + 1)$ vertices formed by putting r copies of J in a circle and making one non-internal vertex in each copy of J adjacent to one non-internal vertex in the next copy.

By Lemma 15, collapsing the internal vertices in a copy of J into one vertex cannot increase the optimal pebbling number. Doing this in each copy of J produces C_{3r} . By Theorem 9, we obtain $\Pi_{OPT}(G) \geq 2r = 2n/(k + 1)$.

For general n , let $r = \lfloor n/(k + 1) \rfloor$. Form J' from $K_{n-(r-1)(k+1)}$ by deleting one edge. Form G' by the construction for G above, using one copy of J' and $r - 1$ copies of J . Collapsing $n + 1 - r(k + 1)$ internal vertices of J' into one vertex turns G' into the example G for $r(k + 1)$ vertices. By Lemma 15, $\Pi_{OPT}(G') \geq \Pi_{OPT}(G) \geq 2r \geq 2(n - k)/(k + 1)$. \square

Corollary 11 shows that the construction in Proposition 16 is extremal for $k = 2$, where it produces C_n . For $k = 3$, it provides connected n -vertex graphs with optimal pebbling number asymptotic to $n/2$; the upper bound from Corollary 11 remains $\lceil 2n/3 \rceil$. As k grows, the coefficient on n in Proposition 16 decreases.

However, for $k > 15$ the optimal pebbling number of our next construction exceeds $2\frac{n}{k+1}$ asymptotically for large n . In particular, there is an n -vertex graph G_n with minimum degree k such that $\Pi_{OPT}(G_n) \frac{k+1}{n} \rightarrow 2.4 - \frac{24}{5k+15}$. This limit exceeds 2 when $k > 15$. We present the construction only for $k \equiv 0 \pmod{3}$; slightly weaker results hold for general k .

We will apply Lemma 15 to a graph that we will contract to a cycle. We first develop a lower bound for 2-solvable distributions on cycles.

Lemma 17. *Let G be a graph with distribution D , and let A be a subset of $V(G)$ such that each vertex in A has a neighbor in A . If each vertex in A is 2-reachable under D , then each vertex in A is 2-reachable under any distribution produced from D by a smoothing move.*

Proof. Let D' be a distribution obtained from D by a smoothing move from v . Note that $D'(v) \geq 1$, by the definition of smoothing. By Lemma 6, every vertex of $A - \{v\}$ is 2-reachable under D' . Hence we may assume that $v \in A$.

Let u be a neighbor of v in A , and let σ be a pebbling sequence under D' after which u has two pebbles. If σ has a move out of v , then truncating σ yields a pebbling sequence

showing that v is 2-reachable. Otherwise, v retains at least one pebble after executing σ , and then a pebbling move from u to v gives it another. \square

Lemma 18. *For $n \geq 3$, if at least $n - 1$ vertices are 2-reachable under a distribution D on C_n , then $|D| \geq n$.*

Proof. Having a 2-reachable vertex requires that D has two pebbles on some vertex. This completes the proof when there is at most one unoccupied vertex. Hence we may choose distinct unoccupied vertices u and v . With $n - 1$ vertices 2-reachable under D , Lemma 17 and the weight argument used in the Smoothing Lemma allow us to assume that D is smooth.

Let P and P' be the u, v -paths along the cycle. Since at least $n - 1$ vertices are 2-reachable, we may assume that u is 2-reachable. By Lemma 8, a pebbling sequence cannot move a pebble out of v without using an edge in both directions, which by the No-Cycle Lemma does not occur in some pebbling sequence that moves two pebbles to u . Lemma 8 also implies that u is not 2-reachable under the restrictions of D to P or P' . Therefore, 2-reachability of u requires moving a pebble to u from each of P and P' , independently. Hence each must have a vertex with two pebbles.

In particular, there is a vertex with two pebbles on each path of occupied vertices joining two unoccupied vertices, and therefore $|D| \geq n$. \square

Let $G_{r,s}$ be a graph formed from r disjoint copies of K_s in a row by making each vertex adjacent to all but one vertex in each neighboring copy of K_s (there is only one such graph, up to isomorphism). Similarly, let $H_{r,s}$ be a graph formed from r disjoint copies of K_s in a circle via the same definition (the isomorphism class is determined by the placement of edges joining the last two copies of K_s).

Theorem 19. *For $s \geq 3$ and $r \geq 1$, $\lceil 4r/5 \rceil \leq \Pi_{OPT}(G_{r,s}) \leq 4 \lceil r/5 \rceil$. If $r, s \geq 3$, then $\lceil 4r/5 \rceil \leq \Pi_{OPT}(H_{r,s}) \leq 4 \lceil r/5 \rceil$. The lower bounds hold also when $s = 2$.*

Proof. Call the initial copies of K_s the “cliques”. By dividing the cliques into consecutive groups of five and placing four pebbles on some vertex in the central clique of the group, we obtain a solvable distribution that uses $4 \lceil r/5 \rceil$ pebbles. In particular, note that if $s \geq 3$, then any two vertices in cliques that are two apart in the ring have a common neighbor in the intervening clique. This fails when $s = 2$, and hence the upper bounds require $s \geq 3$ (the lower bound for $s = 2$ is strengthened in Theorem 28).

Now consider the lower bounds. The proof is by induction on r . When $r \leq 4$, the claims are easily checked. Since adding edges to a graph cannot increase the optimal pebbling number, it suffices in the induction step to prove the lower bound for $H_{r,s}$.

For $r \geq 5$, consider a minimal solvable distribution D on $H_{r,s}$. Label the cliques F_1, F_2, \dots, F_r in order. Let A be the set of cliques containing no vertex that is 2-reachable under D . If $|A| \leq 1$, then collapsing each F_i to a single vertex yields a distribution on C_r under which at least $r - 1$ vertices are 2-reachable. By Lemma 18, in this case $|D| \geq r > \lceil 4r/5 \rceil$.

We may therefore assume that $|A| \geq 2$. Suppose first that $F_i \in A$ but $F_{i-1}, F_{i+1} \notin A$. Let u be a 2-reachable vertex in F_{i-1} , and let v be a 2-reachable vertex in F_{i+1} . Since $|A| \geq 2$, we may also choose $F_j \in A$. Since $F_i, F_j \in A$, we can never put two pebbles on any vertex in $F_i \cup F_j$, and hence we can never move a pebble out of $F_i \cup F_j$. Since u and v are separated by $F_i \cup F_j$, this implies that u and v are 2-reachable simultaneously; that is, the pebbles used in moving two pebbles to one are not used in moving two pebbles to the other. Since $s \geq 3$, u and v have a common neighbor w in F_i . Now w is 2-reachable using pebbles moved from u and v , which contradicts $F_i \in A$.

It follows that for every member of A , some neighboring clique is also in A . When $F_i, F_{i+1} \in A$, we call the edges joining F_i and F_{i+1} *useless*. Since we cannot move two pebbles to any vertex in either clique, we cannot move a pebble along an edge joining them. Hence deleting these edges does not affect the solvability of D .

Since for every member of A there is a neighboring clique also in A , every clique in A is incident to a useless set of edges. Hence there are at least $|A|/2$ such useless sets of edges.

If $|A| \geq 3$, then there are at least two useless sets of edges; deleting them leaves a graph whose components are $G_{t,s}$ and $G_{r-t,s}$, with the distribution D still solvable. Applying the induction hypothesis to the two components yields $|D| \geq (4/5)r$.

Otherwise, $|A| = 2$. Lemma 15 implies that collapsing each clique to a single vertex and collapsing the two vertices arising from A to a single vertex v yields a distribution on C_{r-1} under which every vertex except v is 2-reachable. Since its size is $|D|$, Lemma 18 implies that $|D| \geq r - 1 \geq 4r/5$. \square

Corollary 20. *Let k be a positive multiple of 3. For $n \geq k + 3$, there is an n -vertex graph G with minimum degree k such that $\Pi_{OPT}(G) \geq (2.4 - \frac{24}{5k+15} - o(\frac{1}{n})) \frac{n}{k+1}$. When n is a multiple of $(k/3) + 1$, the term $-o(\frac{1}{n})$ can be dropped.*

Proof. Given such n and k , let $s = k/3 + 1$ and $r = \lfloor n/s \rfloor$. Note that $r, s \geq 3$. The graph $H_{r,s}$ is $3(s - 1)$ -regular, since each vertex has $s - 1$ neighbors in its own clique and in each

neighboring clique. Form G by adding to $H_{r,s}$ a set of $n - rs$ vertices whose neighborhoods duplicate neighborhoods of vertices in $H_{r,s}$. Thus G has n vertices and minimum degree at least k . Also $H_{r,s}$ is obtained from G by collapsing sets of vertices. Thus $\Pi_{OPT}(G) \geq \Pi_{OPT}(H_{r,s})$, by Lemma 15. When n is a multiple of s , we compute

$$\frac{\Pi_{OPT}(G)(k+1)}{n} = \frac{\Pi_{OPT}(H_{r,s})(3s-2)}{rs} \geq \frac{4r}{5} \frac{3s-2}{rs} = \frac{12}{5} - \frac{8}{5s} = \frac{12}{5} - \frac{24}{5k+15}.$$

In general, $n \leq rs + s - 1$, so we replace the rs in the denominator above with $rs(1 + \frac{s-1}{rs})$. Since $(1 + \frac{s-1}{rs})^{-1} \geq 1 - \frac{s-1}{rs} \geq 1 - \frac{k}{3n-k}$, we obtain $\frac{12}{5} - \frac{24}{5k+15} - o(\frac{1}{n})$ as a lower bound. \square

Let $f(k)$ be the infimum of all α such that $\Pi_{OPT}(G) \frac{k+1}{n(G)} \leq \alpha$ for all graphs with minimum degree k . By Corollary 14, Proposition 16, and Corollary 20, $\max\{2, 2.4 - \frac{24}{15k+5}\} \leq f(k) \leq 4$. Given the simplicity of Corollary 14, we believe that $f(k)$ is bounded away from 4, but we have no conjecture for an asymptotic value.

For $k = 3$, the upper bound $\Pi_{OPT}(G) \leq 2n(G)/3$ yields $f(3) \leq 8/3$. We have no construction needing more than the $n(G)/2$ of Proposition 16; Theorem 28 provides another such example.

Question 21. *Is it true that $\Pi_{OPT}(G) \leq \lceil n/2 \rceil$ whenever G is a connected n -vertex graph with minimum degree at least 3? The bound would be sharp for $n \geq 6$.*

When $k = 4$, Corollary 20 does not apply, but more pebbles may be needed than the $2n/5$ in Proposition 16. We base our construction on the ‘‘Sierpinski Triangle’’.

Example 22. Let G_1 be a triangle; its three vertices are its *corners* $\{x, y, z\}$. For $m > 1$, given three copies of G_{m-1} with corner vertices $\{x_i, y_i, z_i\}$ in the i th copy, form G_m by collapsing the pairs $\{z_1, x_2\}$, $\{y_2, z_3\}$, and $\{x_3, y_1\}$. The remaining corner vertices $\{x_1, y_2, z_3\}$ are the corners of G_m . Another way to construct G_m from G_{m-1} , starting with a layout of G_1 in the plane, is to subdivide the edges of each bounded triangle and add a new triangle joining each such set of three new vertices.

For $m > 1$, form H_m from G_m by adding three edges to make the corners pairwise adjacent. Since the corners of G_m have degree 2 and all other vertices of G_m have degree 4, H_m is 4-regular for $m > 1$. Also, $n(H_m) = n(G_m) = 3n(G_{m-1}) - 3$; with $n(G_1) = 3$, we have $n(H_m) = (3^m + 3)/2$.

For $m \geq 3$, we present a solvable distribution on H_m with $2 \cdot 3^{m-2}$ pebbles (there are many such distributions), and we conjecture that this is optimal. If so, then $\Pi_{OPT}(H_m)/n(H_m)$ approaches $4/9$ from below.

In forming G_m , three copies of G_{m-1} are used. Further breakdown shows that 3^{m-3} copies of G_3 are used. The number a_m of vertices of G_m that are corners of copies of G_3 equals $n(G_{m-2})$, by the alternative construction. Since the corners of G_3 form a distance-2 dominating set of G_3 , we have $\Pi_{OPT}(H_m) \leq \Pi_{OPT}(G_m) \leq 4n(G_{m-2}) = 2 \cdot 3^{m-2} + 6$.

For $m \geq 3$, we can save six pebbles in this solvable distribution on H_m . The distance between corners of G_m is 2^{m-1} . In H_m , these corners are pairwise adjacent. Hence the four pebbles on one corner x can satisfy the other corners y and z and the immediate neighbors of y and z . Let P be the shortest y, z -path. If we delete the pebbles on P , then the unreachable vertices are within distance 1 of P . By putting two pebbles each on the corners of copies of G_2 along P , we have deleted $4(2^{m-3} + 1)$ pebbles and added $2(2^{m-2} - 1)$ pebbles, saving 6. \square

6 Girth and Minimum Degree

Forbidding short cycles restricts the input in a way that improves upper bounds on the optimal pebbling number. In particular, if G has minimum degree k and girth at least 5, then four pebbles at a vertex v can take care of $k^2 + 1$ vertices, because the neighborhoods of the neighbors of v overlap only at v .

Proposition 23. *If G is a connected graph with minimum degree k and girth at least $2t + 1$, then $\Pi_{OPT}(G) \leq 2^{2t}n/c_k(t)$, where $c_k(t) = 1 + k \sum_{i=1}^t (k - 1)^{i-1}$.*

Proof. When G has minimum degree k and girth at least $2t + 1$, every distance- t neighborhood has size at least $c_k(t)$. Proposition 13 then applies. \square

Note that $c_k(t) = 1 + [(k - 1)^t - 1](1 + \frac{2}{k-2}) > (k - 1)^t$ for fixed k . For fixed k with $k \geq 6$, this yields $\Pi_{OPT}(G)/n(G) \rightarrow 0$ as $t \rightarrow \infty$. A more detailed analysis improves the upper bound. The idea is to use 2^{2t} pebbles on a vertex of the distance- $2t$ dominating set only when it is used to reach substantially more than the $c_k(t)$ vertices guaranteed in its distance- t neighborhood.

Theorem 24. *Let k and t be positive integers with $k \geq 3$ and $t \geq 2$, except not $(k, t) = (3, 2)$. If G is an n -vertex graph with minimum degree k and girth at least $2t + 1$, then $\Pi_{OPT}(G) \leq 2^{2t}n/(c_k(t) + c'(t))$, where $c_k(t)$ is defined as above and $c'(t) = (2^{2t} - 2^{t+1})\frac{t}{t-1}$.*

Proof. As constructed in the proof of Proposition 13, we begin with a distance- $2t$ dominating set S of size at most $n/c_k(t)$, where $c_k(t)$ is defined as in Proposition 23 and the distance between any two vertices of S is at least $2t + 1$.

To each $v \in S$, we assign a set $R(v)$ of vertices in G ; pebbles on v will be used to reach the vertices of $R(v)$. Each vertex within distance t of v is in $R(v)$; this causes no conflict, since the distance- t neighborhoods from vertices of S are disjoint. Indeed, we grow the sets of the form $R(v)$ to absorb all vertices of G by doing a simultaneous breadth-first search from all of S ; each vertex goes into just one of these sets when it is reached. Since S is a distance- $2t$ dominating set, for each $v \in S$ this generates a spanning tree $T(v)$ of the subgraph induced by $R(v)$, such that leaves of $T(v)$ have distance at most $2t$ from v in $T(v)$.

Let $R'(v)$ be the set of nonleaf vertices of $T(v)$ that are not within distance t of v . Let $r'(v) = |R'(v)|$. If $r'(v) < 2^{2t} - 2^{t+1}$, then put 2^{t+1} pebbles on v and one pebble on each vertex of $R'(v)$. Otherwise, put 2^{2t} pebbles on v .

When $r'(v) \geq 2^{2t} - 2^{t+1}$, the 2^{2t} vertices on v can reach all vertices at distance at most $2t$ from v . When $r'(v) < 2^{2t} - 2^{t+1}$, the 2^{t+1} pebbles on v can reach vertices at distance $t + 1$ from v , including the closest ones in $R'(v)$. The rest of $T(v)$ can then be reached by pebbling along paths through $R'(v)$. Hence the distribution is solvable.

When $r'(v) < 2^{2t} - 2^{t+1}$, we use $r'(v)$ pebbles on $R'(v)$. We claim that at least $r'(v)\frac{t}{t-1}$ vertices lie in $T(v)$ that are not within distance t of v . For $0 \leq i \leq t-1$, let p_i be the number of vertices in $T(v)$ that are i levels above a leaf, but not within distance t of v . For $i > 0$, the vertices counted by p_i have distinct children in $T(v)$ counted by p_{i-1} , so $p_0 \geq p_1 \geq \dots \geq p_{t-1}$. Also, $r'(v) = \sum_{i=1}^{t-1} p_i$. We put pebbles on $r'(v)$ vertices, but we add $r'(v) + p_0$ vertices beyond those counted by $c_k(t)$. We have $\frac{r'(v) + p_0}{r'(v)} = 1 + \frac{p_0}{r'(v)} \geq 1 + \frac{p_0}{(t-1)p_0} = \frac{t}{t-1}$. Hence we add at least $r'(v)t/(t-1)$ vertices not previously counted.

We have shown that when $r'(v) < 2^{2t} - 2^{t+1}$, we use $2^{t+1} + r'(v)$ pebbles with $T(v)$ having at least $c_k(t) + r'(v)\frac{t}{t-1}$ vertices. When $r'(v) \geq 2^{2t} - 2^{t+1}$, we use 2^{2t} pebbles, with $T(v)$ having at least $c_k(t) + c'(t)$ vertices.

Let $S' = \{v \in S : r'(v) < 2^{2t} - 2^{t+1}\}$, and let $s = |S'|$. Let $r = \sum_{v \in S'} (2^{2t} - 2^{t+1} - r'(v))$. We have $n \geq s[c_k(t) + c'(t)] - r\frac{t}{t-1}$, and we used $2^{2t}s - r$ pebbles. Thus

$$\Pi_{OPT}(G) \leq \frac{2^{2t}s - r}{s[c_k(t) + c'(t)] - r\frac{t}{t-1}}n \leq \frac{2^{2t}}{c_k(t) + c'(t)}n,$$

where the last inequality uses that $2^{2t}/(c_k(t) + c'(t)) < (t - 1)/t$ when $k \geq 3$ and $t \geq 2$ and $(k, t) \neq (3, 2)$. \square

Since $c'(t) \geq 2^{2t}$ and $c_4(t) = 1 + (4^t - 1)(5/3)$, the resulting upper bound on $\Pi_{OPT}(G)/n(G)$ when $k = 5$ tends to $3/8$ as $t \rightarrow \infty$. For $k = 2$, always $\Pi_{OPT}(C_n) = \lceil 2n/3 \rceil$. Thus it is natural to ask whether the behavior we noted for $k \geq 6$ also holds for $3 \leq k \leq 5$.

Question 25. For $k \in \{3, 4, 5\}$, does there exist $f_k(t)$ such that $\lim_{t \rightarrow \infty} f_k(t) = 0$ and graphs with minimum degree k and girth at least $2t + 1$ satisfy $\Pi_{OPT}(G)/n(G) \leq f_k(t)$?

We have not constructed graphs to show that the bound in Theorem 24 is sharp, and we do not believe that it is sharp. We present one more result, showing that if G has girth 4 and minimum degree 4, then $\Pi_{OPT}(G)$ can be as large as $n(G)/2$. This improves the construction in Proposition 16 for $k = 3$ by showing that even when triangles are forbidden the same number of pebbles may be needed.

The *cartesian product* $G \square H$ of graphs G and H is the graph with vertex set $V(G) \times V(H)$ such that (u, v) is adjacent to (u', v') if and only if (1) $u = u'$ and $vv' \in E(H)$ or (2) $v = v'$ and $uu' \in E(G)$. Note that $G \square H$ contains a copy of H for each vertex of G and a copy of G for each vertex of H .

In particular, $C_m \square K_2$ and $P_m \square K_2$ are circular and linear “ladders”; two copies of the cycle or path, with corresponding vertices from the two copies adjacent. We call the m copies of K_2 the *rungs* of the graph. In $C_m \square K_2$, exchanging the matching joining two rungs for the other possible matching joining them yields a graph isomorphic to the graph formed from a $2m$ -cycle by adding chords joining opposite vertices (those at distance m along the cycle). This graph has been called the “Möbius ladder”, so we denote it by M_m .

The graphs $C_m \square K_2$ and M_m are special cases of the construction in Theorem 19 with $m = r$ and $s = 2$. The lower bound there is $4m/5$; this result improves that bound. To prove the lower bound, we need to characterize the optimal 2-solvable distributions on paths. For this we need an analogue of Lemma 7 for 2-solvable distributions.

Lemma 26. *Every connected graph with at least three vertices (other than a cycle) has a smooth minimal 2-solvable distribution that gives at most two pebbles to each leaf.*

Proof. We apply smoothing moves to a minimal 2-solvable distribution on such a graph G . Since every vertex is 2-reachable, every vertex has a 2-reachable neighbor, and hence the result of a smoothing move is also a 2-solvable distribution, by Lemma 17. We showed in

the proof of Lemma 7 that when G is not a cycle only finitely many smoothing moves can be made, so we obtain a smooth minimal 2-solvable distribution D .

Suppose now that $D(v) > 2$ for some leaf v . Let u be its neighbor, and let $j = D(u)$ and $k = D(v) \geq 3$. Obtain D' from D by setting $D'(v) = 1$ and $D'(u) = j + k - 2$; leave other values unchanged. Now D' starts with at least as many pebbles on u as v could send there under D to help pebble other vertices. If $j + k \geq 4$, then $D'(u) \geq 2$ to provide a second pebble for v . Otherwise, $(j, k) = (0, 3)$; now v can send only one pebble to u under D , so the 2-solvability of D requires that another pebble can be moved to join the pebble on u under D' ; they can then provide a second pebble for v . Hence D' is 2-solvable, but $|D'| < |D|$, which contradicts the minimality of D . \square

A slightly longer case analysis ensures a smooth 2-solvable distribution with at most one pebble on each leaf, but we will not need this.

Theorem 27. *Every 2-solvable distribution on P_n has at least $n + 1$ pebbles. Furthermore, the 2-solvable distributions with $n + 1$ pebbles consist of “prime segments” separated by single unoccupied vertices, where a prime segment is a path with either (1) two pebbles on one vertex and one pebble on all other vertices, or (2) three consecutive vertices having 0, 4, 0 pebbles, respectively, and one pebble on all other vertices.*

Proof. We use induction on n ; when $n \leq 2$ the unique minimal 2-solvable distributions have $n + 1$ pebbles and are prime segments, as claimed. Consider $n \geq 3$.

By Lemma 26, there is a smooth 2-solvable distribution D having at most two pebbles on each endpoint. By Lemma 8, the endpoints cannot be unoccupied. If every vertex is occupied, then 2-solvability requires some vertex to have two pebbles, and then the minimal distributions have $n + 1$ pebbles and form a single prime segment.

We may therefore assume that some internal vertex v is unoccupied. By Lemma 8, 2-reachability of v requires one pebble to arrive from each side. Since two pebbles cannot arrive at v from one side, pebbles on one side of v cannot be used to obtain 2-solvability of any vertex on the other side. Hence $P_n - v$ consists of two subpaths, each inheriting a 2-solvable distribution (each neighbor of v is 2-reachable using only pebbles on that side, because each can provide a pebble to v). With these paths having l and $n - 1 - l$ vertices, the induction hypothesis requires $l + 1 + n - l$ pebbles in D , and it also completes the decomposition into prime segments after the split at v .

We now consider other optimal 2-solvable distributions on P_n , not necessarily smooth. The transformation in Lemma 26 shows that optimal 2-solvable distributions have at most two pebbles on each leaf, smooth or not. Since smoothing moves preserve 2-solvability but do not discard pebbles, a smoothing move on an optimal 2-solvable distribution will not leave a leaf with at least three pebbles. Hence we can obtain all optimal 2-solvable distributions by “inverting” smoothing moves starting with the distributions we have described.

Such an inversion move changes consecutive pebble values (i, j, k) to $(i - 1, j + 2, k - 1)$, where $i, j, k \geq 1$. Since all values are positive, we can never make an unoccupied vertex occupied by such a move, so the three positions must be within a single original prime segment. We claim that the inversion move maintains the property that 2-solvability within the segment requires pebbles to flow out from the unique vertex with most pebbles on the segment, and pebbles never cross an unoccupied internal vertex. Maintaining these properties, we can never make an inversion move with $j = 1$, because by symmetry we may assume $i = 1$, and the newly unoccupied vertex would not be 2-reachable. Hence the only possible inversion moves change $(1, 2, 1)$ to $(0, 4, 0)$, and there can only be one of these within a prime segment. Segments formed by surrounding $(0, 4, 0)$ with single-pebble vertices are 2-solvable, so this completes the description of the optimal 2-solvable distributions. \square

Theorem 28. $\Pi_{OPT}(C_m \square K_2) = \Pi_{OPT}(P_m \square K_2) = \Pi_{OPT}(M_m) \geq m$ for $m \geq 2$. Equality holds except for $m \in \{2, 5\}$.

Proof. We first provide constructions (except when $m \in \{2, 5\}$) to show that the lower bound is sharp. Observe that three pebbles on one rung can reach all vertices on the two neighboring rungs. Also, four pebbles on two adjacent rungs (two each at opposite corners of the resulting 4-cycle) can reach all vertices on the two neighboring rungs. We can cover the graph with disjoint sets of three or four rungs unless $m \in \{2, 5\}$. For $m = 5$, six pebbles suffice. For $m = 2$, actually $M_2 = K_4$ and two pebbles suffice, but $C_2 \square K_2$ and $P_2 \square K_2$ degenerate to 4-cycles and need a third pebble.

For the lower bound, we use induction on m . For $m = 1$ and $m = 2$, note that $\Pi_{OPT}(P_m \square K_2) = m + 1$. Now consider $m \geq 3$. Since $P_m \square K_2 \subseteq C_m \square K_2$, it suffices to prove the lower bound for $C_m \square K_2$. The argument for $C_m \square K_2$ is valid also for M_m .

Consider an optimal solvable distribution D with $|D| \leq m$; we show that equality holds. If some pebbling sequence from D results in a rung having two pebbles, then collapsing that rung to a vertex yields a graph and distribution under which the resulting vertex is

2-reachable, so we say that the rung is 2-reachable under D . If at least $m - 1$ rungs are 2-reachable, then collapsing each rung to a vertex yields a distribution D' on C_m under which $m - 1$ vertices are 2-reachable. Lemma 18 then yields $|D| = |D'| \geq m$.

Now suppose that at least two rungs R and R' are not 2-reachable under D . The pebbles that arrive in pebbling sequences to reach the two vertices of R arrive from the same direction; otherwise, since no pebble can ever emerge from R' , the two pebbling sequences can be performed independently and R is 2-reachable.

Since both sequences reach R from the same side, and no pebble can emerge from R to the other side (because R is not 2-reachable), D remains solvable on the graph obtained by deleting the edges from R to that rung. If there are two nonadjacent rungs that are not 2-solvable, then doing this for those two rungs splits D into solvable distributions on $P_i \square K_2$ and $P_{m-i} \square K_2$, for some i with $1 \leq i \leq m - 1$. The induction hypothesis applies to both subgraphs, and we obtain $|D| \geq m$.

In the remaining case, there are exactly two rungs R and R' that are not 2-reachable, and they are consecutive. A rung that is not 2-reachable is unoccupied, because if there is one pebble on it, then the sequence to reach the other vertex requires bringing another pebble to the rung. Furthermore, the pebbling sequences that move two pebbles to other rungs cannot use vertices in R or R' , since they are not 2-reachable.

Therefore, deleting R and R' and collapsing the remaining rungs yields a 2-solvable distribution D' on P_{m-2} . If $|D'| \geq m$, we have the desired result. Otherwise, D' is a minimal 2-solvable distribution on P_{m-2} . We use the description of all such distributions, obtained in Theorem 27.

Let S be the rung other than R' that neighbors R ; in the collapsed path, S is an endpoint. Under D' , S can receive two pebbles from its neighbor if the prime segment ends 0, or one pebble from its neighbor to join its original pebble if the segment ends with 1, or no pebbles to join its two original pebbles if the segment ends with 2. In no case can S receive a third pebble. Also, each case leaves no choice in the uncollapsed original distribution D about which vertex of the rung S receives the extra pebble or pair. Without getting a third pebble to S or being able to move two pebbles to either vertex of S , it is not possible under D to reach each vertex of R . \square

References

- [1] N. Alon. Transversal numbers of uniform hypergraphs. *Graphs Combin.* 6(1):1–4, 1990.
- [2] V.I. Arnautov. Estimation of the exterior stability number of a graph by means of the minimal degree of the vertices. *Prikl. Mat. i Programirovanie* 11:3–8, 126, 1974.
- [3] B. Crull, T. Cundiff, P. Feltman, G.H. Hurlbert, L. Pudwell, Z. Szaniszló, Z. Tuza. The cover pebbling number of graphs. *Discrete Math.* 296 (2005), 15–23.
- [4] F.R.K. Chung. Pebbling in hypercubes. *SIAM J. Discrete Mathematics*, 2(4):467–472, 1989.
- [5] A. Czigrinow, private communication.
- [6] T. Friedman and C. Wyels. Optimal pebbling of paths and cycles. *Discrete Mathematics*, submitted (Mathematics ArXiv Article math.CO/0506076).
- [7] G.H. Hurlbert. A survey of graph pebbling. *Congressus Numerantium*, 139:41–64, 1999.
- [8] K. Milans and B. Clark. The complexity of graph pebbling. arXiv’ed in 2005 at <http://arxiv.org/abs/math/0503698>.
- [9] D. Moews. Pebbling graphs. *J. Combin. Theory (B)*, 55:244–252, 1992.
- [10] D. Moews. Optimally pebbling hypercubes and powers. *Discrete Mathematics*, 190(1–3):271–276, August 1998.
- [11] D. Moews. Pebbling graphs. <http://xraysgi.ims.uconn.edu/dmoews/pebbling-graphs.ps>, 2004.
- [12] L. Pachter, H.S. Snevily, and B. Voxman. On pebbling graphs. *Congressus Numerantium*, 107:65–80, 1995.
- [13] C. Payan. Sur le nombre d’absorption d’un graphe simple. *Colloque sur la Théorie des Graphes (Paris, 1974)*, Cahiers Centre Études Recherche Opér. 17, 307–317, 1975.