THE COMPLEXITY OF GRAPH PEBBLING

KEVIN MILANS AND BRYAN CLARK

ABSTRACT. In a graph G whose vertices contain pebbles, a pebbling move uv removes two pebbles from u and adds one pebble to a neighbor v of u. The optimal pebbling number $\hat{\pi}(G)$ is the minimum k such that there exists a distribution of k pebbles to G so that for any target vertex r in G, there is a sequence of pebbling moves which places a pebble on r. The pebbling number $\pi(G)$ is the minimum k such that for all distributions of k pebbles to G and for any target vertex r, there is a sequence of pebbling moves which places a pebble on r.

We explore the computational complexity of computing $\hat{\pi}(G)$ and $\pi(G)$. In particular, we show that deciding whether $\hat{\pi}(G) \leq k$ is NP-complete and deciding whether $\pi(G) \leq k$ is $\Pi_2^{\rm P}$ -complete. Additionally, we provide a characterization of when an unordered set of pebbling moves can be ordered to form a valid sequence of pebbling moves.

1. INTRODUCTION

Let G be a simple, undirected graph and let $p : V(G) \to \mathbb{N} \cup \{0\}$ be a distribution of pebbles to the vertices of G. We refer to the total number of pebbles $\sum_{v} p(v)$ as the size of p, denoted by |p|. A pebbling move uv consists of removing two pebbles from a vertex u with $p(u) \ge 2$ and placing one pebble on a neighbor v of u. After completing a pebbling move uv, we are left with a new distribution of pebbles, which we denote by p_{uv} . Similarly, if $\sigma = u_1v_1, \ldots, u_kv_k$ is a sequence of pebbling moves, denote by p_{σ} the distribution of pebbles that results from making the pebbling moves specified by σ . Although graph pebbling was originally developed to simplify a result in number theory (F.R.K. Chung provides the history [C]), it has since become an object of study in its own right. G.H. Hurlbert presents a detailed survey of early graph pebbling results [H99].

Notational Conventions: We use G and H to refer to simple, undirected graphs. We use D and E to refer to directed graphs with multiple edges and loops. If v is a vertex in a directed multigraph, we denote the indegree (resp. outdegree) of v by $d^+(v)$ (resp. $d^-(v)$). We write n(G) (resp. e(G)) for the number of vertices (resp. edges) in G. Similarly, we use V(G) (resp. E(G)) to refer to the vertex set (resp. edge set) of G. We write $d_G(u, v)$ (or d(u, v) when G is clear from context) for the length of the shortest uv-path in G.

If p and q are pebble distributions on a graph G, we say that $p \ge q$ if $p(v) \ge q(v)$ for each vertex v in G.

Given a graph G with a pebble distribution p, we say that a vertex r in G is reachable if there is a sequence of pebbling moves which places a pebble on r. Note that whenever p(r) > 0, r is trivially reachable. The notion of reachability is fundamental to graph pebbling; most of our decision problems involve questions of reachability. We call the problem of deciding (given G, p, and r) whether r is reachable REACHABLE. The complexity of REACHABLE was first studied by G.H. Hurlbert and H. Kierstead, who found that REACHABLE is NP-complete [HK]. The result was independently rediscovered by N.G. Watson [W]. In section 3, we present another proof that REACHABLE is NP-complete; in later sections, we extend this proof to establish the complexity of OPTIMAL-PEBBLING-NUMBER and PEBBLING-NUMBER.

Given a graph G and a target vertex r, the r-pebbling number of G, denoted $\pi(G, r)$, is the minimum k such that r is reachable under every pebble distribution of size k. Similarly, the pebbling number of G, denoted $\pi(G)$, is the minimum k such that every vertex in G is reachable under every pebble distribution of size k. For a connected graph G, a pigeonhole argument quickly establishes that such a k exists and so $\pi(G)$ is well defined (see Proposition 5.1). We call the problem of deciding whether $\pi(G, r) \leq k$ (resp. $\pi(G) \leq k$) R-PEBBLING-NUMBER (resp. PEBBLING-NUMBER). In section 5, we establish that both decision problems are

Date: 8th June 2005.

Short Name	Full Name	Description	Complexity
P N	PEBBLING-NUMBER	Given G,k : is $\pi(G) \le k$?	$\Pi_2^{\rm P}$ -complete
RPN	R-PEBBLING-NUMBER	Given G, k, r : is $\pi(G, r) \le k$?	$\Pi_2^{\rm P}$ -complete
OPN	OPTIMAL-PEBBLING-NUMBER	Given G, k : is $\widehat{\pi}(G) \leq k$?	NP-complete
PR	REACHABLE	Given G, p, r : is r reachable?	NP-complete
PC	COVERABLE	Given G, p : does p cover the unit distribution?	NP-complete

FIGURE 1.1. A summary of the decision problems considered in this paper

 $\Pi_2^{\rm P}$ -complete, meaning that these problems are complete for the class of problems computable in polynomial time by a coNP machine equipped with an oracle for an NP-complete language. Consequently, these decision problems are both NP-hard and coNP-hard. It follows that R-PEBBLING-NUMBER and PEBBLING-NUMBER are neither in NP nor in coNP unless NP = coNP. N.G. Watson simultaneously and independently established that R-PEBBLING-NUMBER is coNP-hard [W].

Observe that if we fix some r in G and put one pebble on every other vertex, r is not reachable. It follows that $\pi(G) \ge n(G)$. It is natural to wonder which graphs achieve equality in $\pi(G) = n(G)$. Although no characterization of such graphs is known, a growing body of results provide conditions that are necessary or sufficient to imply $\pi(G) = n(G)$. Recall that G is k-connected if $n(G) \ge k + 1$ and for every set S of at most k - 1 vertices, G - S is connected. In [CHH], it is shown that if G is 3-connected and has diameter 2, then $\pi(G) = n(G)$. Consequently, the probability that a random graph on n vertices satisfies $\pi(G) = n(G)$ approaches 1 as n grows. Furthermore, in [CHKT] it is shown that if G has diameter d and is (2^{2d+3}) connected, then $\pi(G) = n(G)$. On the other hand, if G contains a cut vertex, then $\pi(G) > n(G)$. Indeed, suppose v is a cut vertex in G and let u and w be vertices in separate components of G - v. If we put three pebbles on u, zero pebbles on v and w, and one pebble on every other vertex, then it is not possible to place a pebble on w.

The optimal pebbling number of G, denoted $\hat{\pi}(G)$, is the minimum k such that each vertex is reachable under some distribution of size k. We call the problem of deciding whether $\hat{\pi}(G) \leq k$ OPTIMAL-PEBBLING-NUMBER. In section 4, we establish that OPTIMAL-PEBBLING-NUMBER is NP-complete.

It is immediate that $\hat{\pi}(G) \leq n(G)$. However, this bound is not tight for connected graphs. As shown in [BCCMW], if G is connected, then $\hat{\pi}(G) \leq \lceil 2n(G)/3 \rceil$. Equality is achieved by the path [BCCMW, PSV] and the cycle [BCCMW]. It is an open problem to characterize which graphs achieve equality.

Given G and distributions p and q, we say that p covers q if there exists a sequence of pebbling moves σ such that $p_{\sigma} \geq q$. The unit distribution assigns one pebble to each vertex in G. We call the problem of deciding whether p covers the unit distribution COVERABLE. In section 3, we establish that COVERABLE is NP-complete; this result was obtained simultaneously and independently by N.G. Watson [W].

Although most of the problems we study are computationally difficult, there are some interesting pebbling problems that are tractable. A pebble distribution q is *positive* if q assigns at least one pebble to every vertex. A distribution p is *simple* if it assigns zero pebbles to all but one vertex. The q-cover pebbling number of G, denoted $\gamma_q(G)$, is the minimum k such that every distribution of size k covers q. The cover pebbling theorem states that for any positive distribution q, there is a simple distribution p of size $\gamma_q(G) - 1$ such that p does not cover q [VW, S]. As a consequence, given G and a positive distribution q, one easily computes $\gamma_q(G)$ in polynomial time. In the special case that q is the unit distribution, we simply write $\gamma(G)$ for $\gamma_q(G)$.

Overview: In section 2, we develop a characterization of when unordered sets of pebbling moves may ordered in a way that yields a valid sequence of pebbling moves. In section 3, we present results on the complexity of REACHABLE and COVERABLE. We also observe that a simple greedy strategy solves REACHABLE whenever G is a tree. Section 3 uses some results from section 2. In section 4, we present our results on the complexity of computing the optimal pebbling number. Section 4 uses some results from sections 2 and 3. In section 5, we present our results on the complexity of computing the (r)-pebbling number. This section uses some results from sections 2 and 3; it is generally independent of section 4. In section 6, we present our conclusions.

Let us consider a simple example. Suppose we are given a graph H with a distribution of pebbles, and we wish to determine if there is a sequence of pebbling moves which ends with only one pebble left in the entire

graph. We call this problem ANNIHILATION. It is not difficult to see that ANNIHILATION is NP-hard. Indeed, a reduction from HAMILTONIAN-PATH is almost immediate. Specifically, to decide if G has a Hamiltonian path, we may construct H from G by introducing a new vertex v which is adjacent to each vertex in G. We place two pebbles on v and one pebble on every other vertex in H. It is clear that G has a Hamiltonian path if and only if there is a sequence of pebbling moves which results in only one pebble in H.

What is less clear is that ANNIHILATION is in NP. If σ is a sequence of pebbling moves in G under p which results in only one pebble left in G, then the length of σ is |p| - 1, which may be exponentially large in the number of bits needed to represent G and p. Hence, σ may be too large to serve as a certificate for membership in ANNIHILATION. However, as we will see, the order of the moves in σ is insignificant. In fact, if we are merely told how many times σ pebbles along each direction in every edge in G, then we can quickly verify the existence of σ .

2. PEBBLE ORDERABILITY

Many questions in graph pebbling concern the existence of a sequence of pebbling moves with certain properties. There is a natural temptation to search for such sequences directly, by deciding which pebbling move to make first, which to make second, and so forth. In this section, we develop tools that allow us more flexibility in constructing sequences of pebbling moves. In particular, our goal is to worry only about which moves we should make, and not the order in which to make them.

We define the signature of a sequence of pebbling moves σ in a graph G to be the directed multigraph on vertex set V(G) where the multiplicity of an edge uv is the number of times σ pebbles from u to v. We say that a digraph D is orderable under a pebble distribution p if some ordering of E(D) is a valid sequence of pebbling moves. We characterize when D is orderable under p. We call the problem of testing whether D is orderable under p ORDERABLE, or PO for short.

As it turns out, two conditions which are necessary for D to be orderable are also sufficient. Suppose that D is orderable and consider a vertex v. We note that v begins with p(v) pebbles, D pledges that v will receive $d_D^-(v)$ pebbles from pebbling moves into v, and D requests $d_D^+(v)$ pebbling moves out of v. Because each pebbling move out of v costs two pebbles, it is clear that $p(v) + d_D^-(v)$ is at least $2d_D^+(v)$. This leads us to define the *balance* of a vertex v as

balance
$$(D, p, v) = p(v) + d_D^-(v) - 2d_D^+(v).$$

The balance of v is simply the number of pebbles that remain on v after executing any sequence of pebbling moves whose signature is D; that is, for any σ whose signature is D, we have that $p_{\sigma}(v) = \text{balance}(D, p, v)$.

If D is orderable under p, then the balance of each vertex must be nonnegative. We call this condition the balance condition. The balance condition alone is not sufficient: if D is a directed cycle and each vertex has one pebble, then the balance of each vertex is zero but we cannot make any pebbling moves, and so Dis not orderable. However, as was implicitly observed in [M], if D is acyclic, then the balance condition is sufficient.

Theorem 2.1 (Acyclic Orderability Characterization). [M] If D is an acyclic digraph with distribution p, then D is orderable if and only if the balance condition is satisfied.

Proof. We have observed that the balance condition is necessary. Conversely, if the balance condition is satisfied, then we obtain a sequence of pebbling moves σ whose signature is D by iteratively selecting a source u in D, making all pebbling moves out of u, and deleting u from D.

Despite the simplicity of the acyclic orderability characterization, we are already able to obtain one of our most useful corollaries. It makes precise our intuition that if we are trying to place pebbles on a target vertex r, it is never advantageous to pebble around in a cycle. Our proof is somewhat shorter than previous proofs.

Corollary 2.2 (No Cycle Lemma). [CCFHPST, M] Suppose D is orderable under p. There exists an acyclic $D' \subseteq D$ such that D' is orderable and $\text{balance}(D', p, v) \geq \text{balance}(D, p, v)$ for all v.

Proof. Let D' be a digraph obtained by iteratively removing cycles from D until no cycles remain. Observe that removing a cycle C does not change the balance of vertices outside of C and increases the balance of

vertices in C by one. It follows that $\operatorname{balance}(D', p, v) \geq \operatorname{balance}(D, p, v) \geq 0$ for all v. Hence, D' is acyclic and satisfies the balance condition. By the acyclic orderability characterization, D' is orderable.

In most contexts, if a sequence σ of pebbling moves satisfies certain criterion, then so will any sequence σ' provided that $p_{\sigma'} \ge p_{\sigma}$. As we have seen, in these situations, we are able to restrict our attention to sequences of pebbling moves whose signatures are acyclic. Indeed, all of our major results fall into this category and therefore only require the orderability characterization for acyclic digraphs.

Nevertheless, one may wish to study the existence of sequences of pebbling moves which purposefully remove pebbles from the graph, as in the ANNIHILATION decision problem. Let us return to our orderability characterization for arbitrary D. As we have seen, in general the balance condition is not sufficient. However, as we show in our next lemma, a directed cycle with one pebble on each vertex is the only minimal, nontrivial situation which satisfies the balance condition and does not allow us to make any pebbling moves.

Lemma 2.3. Suppose that D with distribution p satisfies the balance condition, D is connected, and $e(D) \ge 1$. If we cannot make any pebbling move described by an edge in D, then D is a directed cycle and each vertex has exactly one pebble.

Proof. Observe that D does not have any source vertices. Indeed, if v were a source, then the balance condition implies that v has enough pebbles to make all pebbling moves out of v requested by D. Therefore v must have outdegree zero, and so v is an isolated, loopless vertex, which contradicts that D is connected and contains an edge.

Let e = E(D), let $X \subseteq V(D)$ be the set of all sinks, let Y = V(D) - X be the set of all nonsinks, let k be the number of edges with sources in Y and sinks in X, and let z be the number of nonsinks that have exactly one pebble. Note that $e = \sum_{v} d^{-}(v) = k + \sum_{v \in Y} d^{-}(v)$ and $e = \sum_{v} d^{+}(v) = \sum_{v \in Y} d^{+}(v)$. Furthermore, for each $v \in Y$, we have that $p(v) \leq 1$; otherwise, $p(v) \geq 2$ and v has outdegree at least one, contradicting that there are no pebbling moves available. It follows that $z = \sum_{v \in Y} p(v)$. Adding the inequality balance $(D, p, v) \geq 0$ over all $v \in Y$, we obtain

$$\sum_{v \in Y} d^{-}(v) + \sum_{v \in Y} p(v) \ge 2 \sum_{v \in Y} d^{+}(v),$$

or equivalently $e - k + z \ge 2e$, and so $e + k \le z$. Because *D* has no sources, every vertex has indegree at least one and so $e \ge n$. Therefore $n \le e \le e + k \le z \le n$. It follows that n = e = z, so that every vertex in *D* is neither a sink nor a source and has exactly one pebble. Furthermore, because e = n, each vertex in *D* has indegree and outdegree exactly one. It follows that *D* is a directed cycle.

Of course, any sequence of pebbling moves leaves a pebble somewhere in the graph; therefore if D contains an edge and D is orderable, then $\operatorname{balance}(D, p, v) \geq 1$ for some vertex v. In fact, a slight generalization of this observation will serve as our second necessary condition. To develop this condition, we first recall the component digraph.

Let D be a directed multigraph. A strongly connected component A is trivial if A consists of a single vertex with indegree and outdegree zero. Define comp(D), the component digraph of D, to be the digraph obtained by contracting each strongly connected component of D to a single vertex.

Suppose that D is orderable, and consider a sink A in comp(D). Because A is a sink component, any pebbling move whose source is in A also has its sink in A; it follows that unless A is trivial, then there must be some vertex v in A with $balance(D, p, v) \ge 1$. We call the condition that every nontrivial sink in comp(D) contains a vertex of positive balance the *sink condition*. Note that in the directed cycle example, each vertex has balance zero, and so it fails the sink condition.

As we now show, the balance condition together with the sink condition are sufficient for D to be orderable. We require a simple proposition.

Proposition 2.4. If D is a strongly connected digraph and D-uv is not strongly connected, then comp(D-uv) contains a single sink A, u is in A, and v is not in A.

Theorem 2.5 (Orderability Characterization). D is orderable under p if and only if

- (1) (balance condition) every vertex has nonnegative balance, and
- (2) (sink condition) every nontrivial sink A in comp(D) contains some vertex with balance at least one

Proof. We have observed that both conditions are necessary. We show that D is orderable under p by induction on e(D). If e(D) = 0, the statement is trivial. In the remaining cases, we assume that D has at least one edge.

We consider the case that there is a source v in D with outdegree at least one. Because balance $(D, p, v) \ge 0$, v has enough pebbles to make all the pebbling moves that D requests out of v. Let σ be an arbitrary ordering of these moves and obtain D' from D by removing all edges whose source is v. We argue that D' is orderable under p_{σ} . It is clear that D' under p_{σ} satisfies the balance condition. Observe that every sink in $\operatorname{comp}(D')$ either consists of v (and is therefore trivial) or is a sink in $\operatorname{comp}(D)$. It follows that every nontrivial sink in $\operatorname{comp}(D')$ is a nontrivial sink in $\operatorname{comp}(D)$ and hence contains some vertex with balance at least one. By induction, D' is orderable under p_{σ} . In the remaining cases, we assume that every source in D is an isolated vertex.

Next, we consider the case where $\operatorname{comp}(D)$ contains a source A with outdegree at least one. Let uv be an edge from a vertex u in A to a vertex v outside of A. We check that A under p satisfies both the balance condition and the sink condition. The balance condition follows from observing that A is a source in $\operatorname{comp}(D)$. Because A is strongly connected and $\operatorname{balance}(A, p, u) \geq 2$, we have that A satisfies the sink condition. By induction, there is an ordering σ of E(A) which is a valid sequence of pebbling moves. We argue that D - E(A) is orderable under p_{σ} . It is clear that D - E(A) under p_{σ} satisfies the balance condition. Because every nontrivial sink in $\operatorname{comp}(D - E(A))$ is a nontrivial sink in $\operatorname{comp}(D)$, D - E(A) satisfies the sink condition. Because every source in D is an isolated vertex, it must be that there is some edge e in D whose sink is u; this edge e is contained in A. Therefore D - E(A) contains fewer edges in D, so that the inductive hypothesis implies that D - E(A) is orderable under p_{σ} . In the remaining cases we assume that every source in $\operatorname{comp}(D)$ is an isolated vertex in $\operatorname{comp}(D)$.

Because $\operatorname{comp}(D)$ is acyclic and every source in $\operatorname{comp}(D)$ is an isolated vertex in $\operatorname{comp}(D)$, it follows that D consists of disjoint, strongly connected components. Because D is orderable if and only if each component of D is orderable, we assume without loss of generality that D is a single, strongly connected component. If we can make a pebbling move uv which leaves D - uv strongly connected, then it is clear that D - uv under p_{uv} satisfies both conditions and so D is orderable.

It remains to consider the case that every possible pebbling move results in a digraph which is no longer strongly connected. By Lemma 2.3, we have that some pebbling move uv is possible.

First, suppose that uv is the only edge out of u. Note that because D is strongly connected, u must have indegree at least one. Furthermore, because uv is a valid pebbling move, we have $p(u) \ge 2$. It follows that $\text{balance}(D, p, u) \ge 1$. It is clear that D - uv under p_{uv} satisfies the balance condition; together with Proposition 2.4, we have that it also satisfies the sink condition. By induction D - uv is orderable under p_{uv} .

Otherwise, let $uw \in E(D)$, $w \neq v$. Let z be a vertex in D with $balance(D, p, z) \geq 1$ (we allow $z \in \{u, w, v\}$), let P be a uz-path, and let Q be a zu-path. Observe that $uv \notin P$ or $uw \notin P$. In the former case, u and z are in the same strongly connected component in D - uv; in the latter case, u and z are in the same strongly connected component with Proposition 2.4 we have that either D - uv under p_{uv} or D - uw under p_{uw} satisfies both conditions. It follows that D is orderable starting from p.

Observe that if D is acyclic, then the sink condition is trivially satisfied, and we recover the acyclic orderability characterization. Our general orderability characterization yields a quick method for checking whether D is orderable, and so ORDERABLE is in P. As a consequence, we see that ANNIHILATION is in NP.

Before we conclude this section, we use our tools to prove some technical lemmas which will be useful in later sections.

Lemma 2.6. Suppose D is acyclic and orderable under p. Then for any vertex w, there exists $D' \subseteq D$ such that D' is orderable and

 $\begin{array}{lll} \mathrm{balance}(D',p,w) & \geq & \mathrm{balance}(D,p,w) + 2d_D^+(w) \\ & \geq & p(w) + d_D^-(w). \end{array}$

Additionally, if $d_D^+(w) > 0$, then we may take D' to be a proper subgraph of D.

Proof. Observe that if uv is an edge in D with v a sink, then D-uv satisfies the balance condition. Let D' be a digraph obtained from D by iteratively deleting edges into sinks other than w until no such edges remain.

Because D' is acyclic and satisfies the balance condition, the acyclic orderability characterization implies that D' is orderable. Observe that w is a sink, or else D' would contain an edge uv with $v \neq w$ a sink. Furthermore, every edge into w in D remains in D'. It follows that balance $(D', p, w) \geq \text{balance}(D, p, w) + 2d_D^+(w)$.

Often, we wish to explore the consequences of the existence of a sequence of pebbling moves with certain properties. In many contexts, considering a minimum sequence of pebbling moves with the properties in question provides us with additional structure. For example, the no cycle lemma implies that a minimum sequence of pebbling moves witnessing that p covers q must be acyclic.

We define a *proper sink* to be a sink with indegree at least one.

Lemma 2.7 (Minimum Signatures Lemma). Let σ be a minimum sequence of pebbling moves in G under p which places at least k pebbles on r with $p(r) \leq k$. If D is the signature of σ , then D is acyclic, contains no proper sinks except possibly r, the outdegree of r is 0, and the indegree of r is k - p(r).

Proof. By the no cycle lemma, D is acyclic, or else we obtain a shorter sequence of pebbling moves placing at least k pebbles on r. By Lemma 2.6, the outdegree of r is zero, or again we obtain a shorter sequence.

Because $d_D^+(r) = 0$, we have $\operatorname{balance}(D, p, r) = p(r) + d_D^-(r)$. Together with $\operatorname{balance}(D, p, r) \ge k$, we have that $d_D^-(r) \ge k - p(r)$. If $d_D^-(r) > k - p(r)$, then $\operatorname{balance}(D, p, r) > k$. Obtain D' from D by deleting one edge into r. Notice that D' satisfies the balance condition and furthermore $\operatorname{balance}(D', p, r) = \operatorname{balance}(D, p, r) - 1 \ge k$. By the acyclic orderability characterization, we obtain a shorter sequence.

If we are interested in minimum sequences of pebbling moves that place k pebbles on some r in a set R of target vertices, the structure of these sequences is further constrained. Not only do their signatures obey the conditions found in the minimum signatures lemma, but the outdegree of each vertex in S is bounded.

Lemma 2.8. Let σ be a sequence of pebbling moves in G under p that places at least k > 0 pebbles on a vertex $r \in R$ which, among all sequences placing at least k pebbles on some vertex in R, minimizes the total number of pebbling moves. Let D be the signature of σ . For each $v \in R$, we have that the outdegree of v is less than k/2.

Proof. Observe that D is acyclic, or else we contradict the no cycle lemma. Suppose for a contradiction that there is $v \in R$ with $d_D^+(v) \ge k/2$. Because k > 0, we have $d_D^+(v) > 0$ and so Lemma 2.6 yields a shorter sequence of pebbling moves placing at least k pebbles on v, a contradiction.

3. PEBBLE REACHABILITY

Recall that the pebbling number of a graph $\pi(G)$ is the minimum k such that every vertex is reachable under every distribution of size k. It is natural, then, to explore the decision problem that results when we fix a particular distribution and target vertex; that is, given G, p, and r, is r reachable? We call this problem REACHABLE, or PR for short. As we show, PR is NP-complete, even when the inputs are restricted so that G is bipartite, has maximum degree three, and each vertex starts with at most two pebbles.

Analogously, fixing the distribution in the cover pebbling number $\gamma(G)$ yields another decision problem: given G and p, does p cover the unit distribution? We call this problem COVERABLE, abbreviated PC. Although deciding whether $\gamma(G) \leq k$ is possible in polynomial time [VW, S], PC is NP-complete.

A sequence of pebbling moves σ is *nonrepetitive* if for every (unordered) pair of vertices $\{u, v\}$, σ contains at most one pebbling move between the vertices u and v. Similarly to PR, we may ask, given G, p, and r, whether r is reachable via a nonrepetitive sequence of pebbling moves. We call this language NPR (nonrepetitive pebble reachability). We show that NPR is NP-complete. Our reduction is from a restricted form of 3SAT whose instances ϕ are all in a canonical form.

Definition 3.1. A 3CNF formula ϕ is in *canonical form* if

- (1) ϕ has at least 2 clauses,
- (2) each clause contains 2 or 3 variables,
- (3) each variable appears at most 3 times in ϕ ,
- (4) each variable appears either once or twice in its positive form, and
- (5) each variable appears exactly once in its negative form



FIGURE 3.1. If $\phi = (w \lor x) \land (w \lor \overline{x}) \land (\overline{w} \lor y \lor z) \land (x \lor \overline{y} \lor \overline{z})$, then $G^{\text{NPR}}(\phi)$ appears above.

It is well known that 3SAT remains NP-complete when (1-3) are required. Suppose ϕ is a 3SAT formula which satisfies (1-3) but not necessarily (4) or (5). Indeed, if a variable x always appears in its positive (negative) form in ϕ , we obtain a simpler, equivalent formula by setting x to true (false), thus removing all clauses containing $x(\bar{x})$. If x appears twice in its negative form, we simply switch all negative occurrences of x to positive occurrences and all positive occurrences of x to negative occurrences. In this way, we obtain an equivalent formula satisfying all of the above. We define R3SAT to be this restricted form of 3SAT.

Our reduction from R3SAT to NPR employs several simple gadgets. The AND gadget is a vertex v that has two input edges and one output edge; initially, v is given zero pebbles. Notice that if σ is nonrepetitive and contains a pebbling move from v along the output edge, then σ must contain pebbling moves into valong both input edges. The OR gadget is identical, except that v is initially given a single pebble. In this case, if σ is nonrepetitive and contains a pebbling move from v along the output edge, then σ must contain a pebbling move into v along one if the input edges. Using 2-ary AND (OR) gadgets, one easily constructs k-ary AND (OR) gadgets.

The variable gadget is a path $v_1v_2v_3$ of length three. The endpoint vertices $\{v_1, v_3\}$ are initially given two pebbles, and the internal vertex v_2 is initially given zero pebbles. The endpoint vertices correspond to the positive occurrence(s) of the variable in ϕ , and the internal vertex corresponds to the negative occurrence of the variable in ϕ . The variable gadget has two or three output edges, depending upon how many times the corresponding variable appears in ϕ . If x_i appears three times in ϕ , then its associated variable gadget X_i has three output edges, one incident to each v_i . If x_i appears twice in ϕ , then X_i has two output edges, one incident to each of v_1 and v_2 . We say that the output edges incident to v_1 and v_3 are positive output edges and the output edge incident to v_2 is the negative output edge.

Given an instance ϕ of R3SAT, we construct $G = G^{\text{NPR}}(\phi)$ as follows. For each variable x_i in ϕ , we introduce a variable gadget X_i in G. For each clause c_j containing $k \in \{2, 3\}$ variables, we introduce a k-ary OR gadget C_j . The output edges of the X_i are identified with the input edges of the C_j in the natural way: if x_i appears in c_j , a positive output edge of X_i is identified with an input edge of C_j , and if $\overline{x_i}$ appears in c_j , the negative output edge of X_i is identified with an input edge of C_j . The output edges of the C_j are connected to the input edges of an m-ary AND gadget A, where m is the number of clauses in ϕ . Finally, the output edge of A is connected to the target vertex r.

Example. If $\phi = (w \lor x) \land (w \lor \overline{x}) \land (\overline{w} \lor y \lor z) \land (x \lor \overline{y} \lor \overline{z})$, then $G^{\text{NPR}}(\phi)$ appears in Figure 3.1 on page 7. **Proposition 3.2.** Let ϕ be an instance of R3SAT with n variables and m clauses. Then $G^{\text{NPR}}(\phi)$ has at most O(n+m) vertices. **Theorem 3.3.** NPR is NP-complete, even when G has maximum degree three and each vertex starts with at most two pebbles.

Proof. It is immediate that NPR is in NP. Let ϕ be an instance of R3SAT and let $G = G^{\text{NPR}}(\phi)$. Observe that each vertex in G starts with at most two pebbles and the maximum degree in G is three.

We claim that ϕ is satisfiable if and only if there is a nonrepetitive sequence of pebbling moves which ends with a pebble on r. Suppose that ϕ is satisfiable via $f : \{x_1, \ldots, x_n\} \rightarrow \{\text{true}, \text{false}\}$. We construct a nonrepetitive sequence of pebbling moves which ends with a pebble on r as follows. For each variable x_i with $f(x_i) = \text{false}$, we make a pebbling move from each endpoint of X_i to the interior vertex of X_i . Notice that after executing these pebbling moves, for each x_i with $f(x_i) = \text{true}$, we have two pebbles on each endpoint of X_i and for each x_i with $f(x_i) = \text{false}$, we have two pebbles on the interior vertex of X_i . Because f satisfies ϕ , each clause gadget C_i has some input edge which is incident to a vertex in a variable gadget with two pebbles. By construction, each vertex in a variable gadget is incident to at most one clause gadget input edge; therefore we are able to make pebbling moves into each clause gadget C_i . By the construction of our clause gadgets, we are then able to make pebbling moves out of each clause gadget and, by construction, along each of the inputs to the *m*-ary AND gadget. It follows that we are able to make a pebbling move along the output of our AND gadget, which places a pebble on r. It is easily observed that our sequence of pebbling moves is nonrepetitive.

Conversely, suppose that σ is nonrepetitive sequence of pebbling moves which ends with a pebble on r. We construct a satisfying assignment f as follows. Because σ contains a pebbling move across the output of the AND gadget A, it follows that σ contains pebbling moves across the output of each clause gadget C_i . Hence, for each clause gadget C_i , σ contains a pebbling move across an input edge e_i of C_i . If e_i is incident to an endpoint of X_j , then we set $f(x_j) = \text{true}$; otherwise, if e_i is incident to the interior vertex of X_j , we set $f(x_j) = \text{false}$. We claim that we do not attempt to set both $f(x_j) = \text{true and } f(x_j) = \text{false}$. Indeed, if we set $f(x_j) = \text{false}$, then σ contains a pebbling move out of the interior vertex v of X_j along an input edge to some clause gadget. Because σ is nonrepetitive, v starts with zero pebbles, and v has degree three, it must be that σ contains pebbling moves from each of the endpoints in X_j into v. Because each endpoint of X_j starts with only two pebbles and σ is nonrepetitive, the moves into v are the only pebbling moves which originate from the endpoints of X_j . Therefore σ does not contain a pebbling move out of an endpoint of X_j along an input edge of a clause gadget, and hence we never attempt to set $f(x_j) = \text{true}$. If the truth values for any variables remain unset, we set them arbitrarily. Now f witnesses that ϕ is satisfiable.

One of the major tools available to us when designing interesting graph pebbling problems is the path; on a path, the pebbling moves available to us are rather limited. If we are in a situation where we need not concern ourselves with pebbling in cycles, then our options on a path become even more limited. Furthermore, if the path is long, it may be difficult to pebble across. Before using paths to reduce NPR to PR, we explore some basic properties.

Lemma 3.4. Let G be a graph which contains an induced path $P = v_0 \dots v_{n+1}$ containing n+2 vertices, and suppose that each of the n internal vertices in P contains c pebbles. Let D be an acyclic signature of a sequence of pebbling moves so that the edge v_1v_0 has multiplicity $a_0 \ge c$. Then the multiplicity of $v_{n+1}v_n$ is at least $2^n(a_0 - c) + c$.

Proof. Observe that the claim is trivial if $a_0 = 0$; we assume that $a_0 \ge 1$. For $1 \le i \le n$, let a_i be the multiplicity of $v_{i+1}v_i$. We claim that for all $1 \le i \le n$, we have that

- (1) $a_i + c \ge 2a_{i-1}$, and
- (2) $a_i \ge a_0$.

Suppose for a contradiction that $i \ge 1$ is the least integer for which (1) or (2) fails, and consider the vertex v_i . By our selection of i, $a_{i-1} \ge a_0$ and therefore D requests at least a_0 pebbling moves out of v_i along edge $v_i v_{i-1}$. Because $a_0 \ge c$ and $a_0 \ge 1$, we have that $2a_0 > c$; hence, by the balance condition at v_i , the indegree of v_i in D is at least one. Because D is acyclic, D contains no edges of the form $v_{i-1}v_i$. Because v_i is an internal vertex in an induced path in G, the only other edge incident to v_i is $v_i v_{i+1}$. It follows that the indegree of v_i in D is exactly the multiplicity of $v_{i+1}v_i$, and so the indegree of v_i in D is a_i . Therefore the balance condition at v_i implies that $a_i + c \ge 2a_{i-1}$, which together with $a_0 \ge c$ and $a_{i-1} \ge a_0$, implies $a_i \ge a_0$.

Solving our recurrence in (1), we find that $a_i \ge 2^i(a_0 - c) + c$.

We use our path lemma to argue that if we can pebble across a long path several times, than we can place many pebbles on the originating endpoint of the path. Together with Lemma 2.6, we obtain the following corollary.

Corollary 3.5. Under the assumptions of Lemma 3.4, there exists $D' \subseteq D$ such that D' is orderable and balance $(D', p, v_{n+1}) \geq 2^{n+1}(a_0 - c) + 2c$. If in addition we have $d_D^+(v_{n+1}) > 0$, then we may take D' to be a proper subgraph of D.

Our reduction used the notion of nonrepetitive sequences of pebbling moves. In fact, there is a natural correspondence between the nonrepetitive sequences of pebbling moves in a graph G and (arbitrary) sequences of pebbling moves in another graph $\mathcal{S}(G, \alpha)$.

Definition 3.6. We obtain $\mathcal{S}(G, \alpha)$ from G by replacing each edge in G with a path containing α internal vertices, so that $d_{\mathcal{S}(G,\alpha)}(u,v) = (1+\alpha) d_G(u,v)$ for any u, v in G. We call these paths one use paths.

As our next lemma shows, the correspondence holds whenever α is sufficiently large with respect to the number of pebbles in G.

Lemma 3.7. Fix a graph G and a parameter $t \ge 0$. Suppose that $\alpha \ge \max\{\lg 2t, 4\lg e(G)\}$ and let $H = S(G, \alpha)$. Let p be a pebble distribution on G of size at most t and define a pebble distribution q on H so that q and p agree on V(G) and q assigns one pebble each to the internal vertices of H's one use paths. We have the following claims.

- (1) If σ is a nonrepetitive sequence of pebbling moves in G, then there exists a sequence of pebbling moves σ' in H such that p_{σ} and $q_{\sigma'}$ agree on V(G).
- (2) Conversely, if σ is a sequence of pebbling moves in H, then there exists a nonrepetitive sequence of pebbling moves σ' in G such that $p_{\sigma'}(v) \ge q_{\sigma}(v)$ for all v in G.

Proof. Claim 1 is clear. Suppose that σ is a sequence of pebbling moves in H. By the no cycle lemma, we may assume without loss of generality that the signature D of σ is acyclic. We define a digraph D' with vertex set V(G) as follows. Let uv be an edge in G and let $u = w_0 \dots w_{\alpha+1} = v$ be the corresponding one use path in H. The multiplicity of the edge uv in D' is the multiplicity of the edge $w_{\alpha}w_{\alpha+1}$ in D. Because D is acyclic, the balance condition implies that if D contains the edge $w_{\alpha}w_{\alpha+1}$, then D contains all edges $w_k w_{k+1}$. It follows that D' is also acyclic. It is easily seen that balance $(D', p, v) \geq$ balance (D, q, v) for each v in D'. By the acyclic orderability characterization, we obtain a sequence of pebbling moves σ' such that $p_{\sigma'}(v) \geq q_{\sigma}(v)$ for all v in G. It remains to show that D' has no edges of multiplicity at least two, so that σ' is necessarily nonrepetitive.

Suppose for a contradiction that uv is an edge in D' with multiplicity at least two; again, let $u = w_0 \dots w_{\alpha+1} = v$ be the corresponding one use path in H. It follows that $w_{\alpha}w_{\alpha+1}$ has multiplicity at least two in D. Recalling that q assigns each of the internal vertices w_i one pebble, Lemma 3.4 implies that the multiplicity of w_0w_1 is at least $2^{\alpha} + 1$. Because each pebbling move reduces the total number of pebbles by one, certainly the size of q is at least $2^{\alpha} + 2$. But $|q| = |p| + \alpha e(G)$ and together with $t \leq 2^{\alpha-1}$ and $\alpha e(G) \leq 2^{\alpha-1}$, we obtain a contradiction.

Corollary 3.8. REACHABLE is NP-complete, even when G is bipartite, has maximum degree three, and each vertex starts with at most two pebbles.

Proof. By the no cycle lemma and the acyclic orderability characterization, PR is in NP. We reduce NPR to PR as follows. Consider a graph G with maximum degree three, a distribution of pebbles p which places at most two pebbles on each vertex in G, and a target vertex r. Let α be the least odd number larger than max $\{\lg 2 | p |, 4 \lg e(G)\}$. Our reduction outputs $H = S(G, \alpha)$ with pebble distribution q as in Lemma 3.7 and target vertex r. Observe that H is bipartite, has maximum degree three, and each vertex starts with at most two pebbles. By Lemma 3.7, r is reachable via a nonrepetitive sequence of pebbling moves in G if and only if r is reachable in H.

Let ϕ be an instance of R3SAT. We define $G^{PR}(\phi) = \mathcal{S}(G^{NPR}(\phi), \alpha)$ with α chosen as in our corollary; that is, G^{PR} is the composition of our reduction from R3SAT to NPR and our reduction from NPR to PR.

Corollary 3.9. COVERABLE is NP-complete, even when G is bipartite, has maximum degree three, and each vertex starts with at most three pebbles.

Proof. By the no cycle lemma and the acyclic orderability characterization, we have that PC is in NP. We reduce PR to PC as follows. Let G be a graph with pebble distribution p and target vertex r. Define a new distribution q of pebbles so that q(v) = p(v) + 1 for all $v \neq r$ and q(r) = p(r). We claim that r is reachable under p if and only if q covers the unit distribution. The forward direction is clear.

Suppose that σ is a minimum sequence of pebbling moves witnessing that q covers the unit distribution, and let D be the signature of σ . By the no cycle lemma, D is acyclic. Because balance $(D, q, v) \ge 1$, we have that balance $(D, p, v) \ge 0$ for all v and balance $(D, p, r) \ge 1$. It follows from the acyclic orderability characterization that D is orderable under p. Together with balance $(D, p, r) \ge 1$, we have that r is reachable under p.

As we have seen, REACHABLE is NP-complete, even under some restrictions of the inputs. However, as we now observe, if we restrict G to be a tree, then we can solve REACHABLE in polynomial time using a simple greedy strategy. A greedy pebbling move is a pebbling move uv such that d(v,r) < d(u,r). The greedy pebbling strategy arbitrarily makes greedy pebbling moves until no greedy pebbling move is possible.

Proposition 3.10 (Greedy Tree Lemma). In a tree T with target r, the maximum number of pebbles that can be placed on r is achieved with the greedy pebbling strategy.

Proof. Suppose for a contradiction that under p, it is possible to place k pebbles on r, but if we make the greedy pebbling move uv, it is no longer possible to place at least k pebbles on r. Let σ be a minimum sequence of pebbling moves placing k pebbles on r, and let D be the signature of σ . By the no cycle lemma, D is acyclic. If D contains the edge uv, then the acyclic orderability characterization implies that D - uv is orderable under p_{uv} , implying it is possible to place k pebbles on r even after pebbling uv. Otherwise, if D does not contain the edge uv, then $d^+(u) = 0$, or else D contains a proper sink other than r, contradicting the minimal signatures lemma. Therefore σ does not contain any pebbling moves out of u, and so uv followed by σ is a legal sequence of pebbling moves placing at least k pebbles on r.

4. Complexity of Optimal Pebbling Number

Recall that the optimal pebbling number $\hat{\pi}(G)$ of a graph G is the least number k such that every vertex is reachable under some distribution of size k. We define OPTIMAL-PEBBLING-NUMBER (abbreviated OPN) to be the problem of deciding, given G and k, whether $\hat{\pi}(G) \leq k$. In this section, we show that OPN is NPcomplete. We observe that OPN is in NP; indeed, we may witness that $\hat{\pi}(G) \leq k$ by providing a distribution p of size k and, for each r, the signature D_r of a sequence of pebbling moves showing that r is reachable. More care is needed to establish that OPN is NP-hard. As in our proof that PR is NP-hard, we establish that OPN is NP-hard through an intermediate decision problem.

Let G be a graph and p be a distribution of pebbles to G. A vertex r is *determinative* if r is reachable under p implies that every vertex in G is reachable under p. Informally, if r is determinative, then no vertices in G are more difficult to pebble than r. Our intermediate decision problem is REACHABLE with the added restriction that r is determinative. We call this problem DPR (determinative pebble reachability).

Proposition 4.1. DPR is NP-complete, even when each vertex starts with at most two pebbles.

Proof. Because REACHABLE is in NP, it is immediate that DPR is in NP as well. We show that our reduction G^{PR} from R3SAT to PR actually produces an instance of DPR. Let ϕ be an instance of R3SAT, and let $G = G^{PR}(\phi)$ with distribution p and target r. We show that r is determinative. Suppose that it is possible to place a pebble on r, or equivalently that ϕ is satisfiable. Consider a vertex $v \in G$. If v is an internal vertex in a one use path introduced in our reduction from NPR to PR, then v begins with one pebble and so v is reachable trivially.

It remains to consider the case that v is a vertex introduced in our reduction from R3SAT to NPR, so that v is either in an OR gadget, a variable gadget, an AND gadget, or v = r. If v is in an OR gadget, then v begins with a pebble. If v is an endpoint of a variable gadget, then v begins with two pebbles. If v is the interior vertex of a variable gadget, then we may place a pebble on v by pebbling from either of the endpoints (which start with two pebbles) across the one use path. Otherwise, if v is in an AND gadget or v = r, then we use the satisfiability of ϕ to place a pebble on v.



FIGURE 4.1. star(3,5)

Before we are able to present our reduction from DPR to OPN, we require some technical lemmas. The following weighting argument is well known and is a fundamental tool in graph pebbling.

Proposition 4.2 (Standard Weight Equation). Let G be a graph with distribution p and target vertex r, and let a_i be the number of pebbles at distance i from r. If it is possible to place s pebbles on r, then we have $\sum_{i>0} 2^{-i}a_i \geq s$.

Proof. Observe that it is not possible to make a pebbling move which increases the sum $\sum_{i>0} 2^{-i}a_i$. \Box

The following graph will be useful to us in two different contexts: first, as a gadget, and secondly in establishing the correctness of our reduction from DPR to OPN.

Definition 4.3. We define $\operatorname{star}(\alpha, \beta)$ to be the result of replacing each edge in $K_{1,\beta}$ with a path of length α , so that $\operatorname{star}(\alpha, \beta)$ has $\alpha\beta$ edges. Equivalently, $\operatorname{star}(\alpha, \beta) = \mathcal{S}(K_{1,\beta}, \alpha - 1)$.

Example. star(3,5) appears in Figure 4.1 on page 11.

Our reduction from DPR to OPN produces a graph whose global structure is similar to that of $\operatorname{star}(\cdot)$. Our instance of DPR plays the role of the center vertex, and the gadgets that we add play the role of the leaves. When we argue the correctness of our reduction, we apply the following lemma to limit the pebble distributions that we must consider. The lemma shows how, despite the simplicity of the standard weight equation, it yields nontrivial results.

Lemma 4.4. Fix $\alpha \geq 1$ and $\beta \geq 2$. Let p be a distribution of $\beta 2^{\alpha}$ pebbles to $\operatorname{star}(\alpha, \beta)$ with the property that for each leaf l in $\operatorname{star}(\alpha, \beta)$, it is possible to place 2^{α} pebbles on l. If $2(\beta^2 + 1) < 2^{\alpha}$, then p is the distribution which places 2^{α} pebbles on each leaf and zero pebbles on the other vertices.

Proof. Let v be the center vertex of $\operatorname{star}(\alpha, \beta)$, and for each $0 \leq i \leq \alpha$, let a_i be the number of pebbles at distance i from v. For each l, it is possible to place 2^{α} pebbles on l and Proposition 4.2 yields an equation; we sum these equations. Because there are β leaves, we obtain $\beta 2^{\alpha}$ on the right hand side. A pebble at distance i from v is at distance $\alpha - i$ from its closest leaf and $\alpha + i$ from all other leaves. It follows that pebbles at distance i from v contribute $1/2^{\alpha-i} + (\beta - 1)/2^{\alpha+i}$ to the left hand side of the equation. We obtain

$$\sum_{i=0}^{\alpha} \left(\frac{1}{2^{\alpha-i}} + \frac{\beta-1}{2^{\alpha+i}} \right) a_i \ge \beta 2^{\alpha}$$

and after some simplification,

$$\sum_{i=0}^{\alpha} \left(2^i + \frac{\beta - 1}{2^i} \right) a_i \ge \beta 4^{\alpha}.$$

Let $f(x) = 2^x + (\beta - 1)2^{-x}$, so that pebbles at distance *i* contribute f(i) to the left hand side. Analyzing the derivative $f'(x) = \ln 2 (2^x - (\beta - 1)2^{-x})$, we find that f'(x) = 0 has one solution, namely $x_0 = \log_4(\beta - 1)$. Furthermore, for $x > x_0$, we have f'(x) > 0 and for $x < x_0$, we have f'(x) < 0. It follows that f(x) has a global minimum at $x = x_0$, f(x) is decreasing on $(-\infty, x_0]$, and f(x) is increasing on $[x_0, \infty)$.

Let $m = \sum_{i=0}^{\alpha-1} a_i$ be the number of pebbles not at distance α from v; we show that m < 1, implying that m = 0. Noting that $a_{\alpha} = \beta 2^{\alpha} - m$, we have that

$$\left(\max_{0 \le i \le \alpha - 1} f(i)\right) m + f(\alpha) \left(\beta 2^{\alpha} - m\right) \ge \beta 4^{\alpha}.$$

Because of the monotonicity properties of f, we have $\max_{i=0}^{\alpha-1} f(i) \in \{f(0), f(\alpha-1)\}$. Because $2(\beta^2+1) < 2^{\alpha}$, certainly $2\beta < 2^{\alpha}$ and therefore

$$f(0) = \beta < 2^{\alpha - 1} \le 2^{\alpha - 1} + (\beta - 1)2^{1 - \alpha} = f(\alpha - 1)$$

It follows that $\max_{0 \le i \le \alpha - 1} f(i) = f(\alpha - 1)$, and after substitution and further simplification, we obtain

$$m \le \frac{\beta 2^{\alpha} \left(f(\alpha) - 2^{\alpha}\right)}{f(\alpha) - f(\alpha - 1)}.$$

Substituting our formula for $f(\alpha)$ into the numerator yields

$$m \le \frac{\beta(\beta - 1)}{f(\alpha) - f(\alpha - 1)} \le \frac{\beta^2}{f(\alpha) - f(\alpha - 1)}.$$

Recall that $2(\beta^2+1) < 2^{\alpha}$, which implies that $\beta^2 < 2^{\alpha-1}-1$. Observe that $f(\alpha) - f(\alpha-1) = 2^{\alpha-1} - (\beta-1)/2^{\alpha}$. Because $\beta - 1 < 2(\beta^2+1) < 2^{\alpha}$, we have that $f(\alpha) - f(\alpha-1) > 2^{\alpha-1} - 1$. It follows that $\beta^2 < f(\alpha) - f(\alpha-1)$, implying that m < 1 as required.

It follows that p places every pebble at distance α from v. It remains to show that p places 2^{α} pebbles on each leaf. Fix an arbitrary leaf l, and let n be the number of pebbles that p places on l. Applying the standard weight equation to l, we have that

$$n + \frac{\beta 2^{\alpha} - n}{2^{2\alpha}} \ge 2^{\alpha}.$$

After simplification, we obtain that

$$n \geq 2^{\alpha} - \frac{2^{\alpha}(\beta - 1)}{4^{\alpha} - 1}.$$

Similarly to the previous paragraph, we show that $n > 2^{\alpha} - 1$. We have that

$$\frac{2^{\alpha}(\beta - 1)}{4^{\alpha} - 1} \le \frac{2^{\alpha}(\beta - 1)}{4^{\alpha} - 2^{\alpha}} = \frac{\beta - 1}{2^{\alpha} - 1}$$

Because $\beta \leq 2(\beta^2 + 1) < 2^{\alpha}$, we have that $(\beta - 1)/(2^{\alpha} - 1) < 1$ and hence $n > 2^{\alpha} - 1$ as required. Therefore p assigns each leaf at least 2^{α} pebbles and the lemma follows.

We now have the tools necessary to present our reduction from DPR to OPN. Let G be a graph with pebble distribution p and determinative target vertex r. Let m = |p|, let $\alpha = \lfloor lg(2(m^2 + 1)) \rfloor$, and let $\beta = 2^{\alpha}m + 2$. We construct a graph H with the property that $\hat{\pi}(H) \leq m2^{\alpha}$ if and only if r is reachable in G.

We construct H from G by attaching a copy of $\operatorname{star}(\alpha, \beta)$ to each pebble in G. That is, for each pebble on a vertex u, we introduce a copy of $\operatorname{star}(\alpha, \beta)$ and attach it to u by identifying u with one of the leaves of our copy of $\operatorname{star}(\alpha, \beta)$.

Lemma 4.5. r is reachable in G under p if and only if $\widehat{\pi}(H) \leq m2^{\alpha}$.

Proof. (\implies). Suppose r is reachable. Define a distribution q of $m2^{\alpha}$ pebbles to H by placing 2^{α} pebbles at the centers of each of the m copies of $\operatorname{star}(\alpha, \beta)$ in H. Consider a vertex v in H. If v belongs to a copy S of $\operatorname{star}(\alpha, \beta)$, then v is at a distance of at most α from the center of S; because the center of S begins with 2^{α} pebbles, v is reachable. Otherwise, v must be a vertex in G. Because r is reachable and determinative under p, to show that v is reachable, it suffices to show that q covers p. But each star can contribute one pebble to the vertex it shares with G, and so q covers p.

(\Leftarrow). Let q be a distribution of $m2^{\alpha}$ pebbles to H witnessing that $\hat{\pi}(H) \leq m2^{\alpha}$. We claim that if u is the center vertex of a copy S of star (α, β) , then it is possible to place 2^{α} pebbles on u starting from q. Indeed, because S contains $\beta - 1 > m2^{\alpha}$ pendant paths with endpoint u, there is some path to which q assigns no pebbles (except possibly at u). Let $w_0w_1\ldots w_{\alpha}$ be one such path with $w_0 = u$. Because every vertex is reachable under q, certainly w_{α} is reachable; let D be a signature of a minimum sequence of pebbling moves

that places a pebble on w_{α} . Because w_{α} is a leaf and q assigns no pebbles to w_{α} , $w_{\alpha-1}w_{\alpha}$ is an edge in D; therefore Corollary 3.5 implies that we can place 2^{α} pebbles on u.

When a graph has a pebble distribution, contracting a set of vertices S changes the pebble distribution in the natural way: pebbles on vertices in S are collected at the vertex of contraction. Construct H' and pebbling distribution q' from H and q by iteratively applying the following contractions:

- (1) Contract all vertices in H that are also in G to a single vertex v
- (2) For each copy S of star(α, β), contract the vertices in S that are at distance at least α from v

Observe that H' is exactly $\operatorname{star}(\alpha, m)$, with center vertex v. Because the contraction operation cannot make pebbling more difficult, it is possible to place 2^{α} pebbles on each leaf in H' starting from q'. Because $2(m^2 + 1) \leq 2^{\alpha}$, applying Lemma 4.4 to $H' = \operatorname{star}(\alpha, m)$ implies that q' must assign 2^{α} pebbles to each leaf of H'. It follows that q assigns 2^{α} pebbles to each copy of $\operatorname{star}(\alpha, \beta)$ in H in such a way that each pebble is at distance at least α away from the vertices in G.

Let E be the signature of a minimum sequence of pebbling moves in H starting from q which places a pebble on r. Consider a copy S of star (α, β) attached to a vertex u in G. We claim that E contains at most one edge from S into u. Indeed, if this were otherwise, then by Corollary 3.5 there exists $E' \subseteq E$ which places at least $2 \cdot 2^{\alpha}$ pebbles on the center vertex of S. However, this is impossible because E is acyclic with edges from S into G and q assigns only 2^{α} pebbles to S.

Obtain E' from E by deleting all edges except those in G. Because each vertex u in G receives a pebble from p for every attached copy of $\operatorname{star}(\alpha,\beta)$, we have that $\operatorname{balance}(E',p,u) \geq \operatorname{balance}(E,q,u)$. It follows from the acyclic orderability characterization that E' is orderable under p; together with $\operatorname{balance}(E',p,u) \geq$ $\operatorname{balance}(E,q,u) \geq 1$, we have that r is reachable under p.

We conclude with this section's main theorem.

Theorem 4.6. OPTIMAL-PEBBLING-NUMBER is NP-complete.

Proof. We have already observed that OPN is in NP and exhibited a reduction from DPR to OPN. It remains to check that the size of H and $m2^{\alpha}$ are not too large, so that our reduction is computable in polynomial time. Let n be the number of vertices in G. By Proposition 4.1, we assume that |p| = m is at most 2n. Our reduction uses gadgets $\operatorname{star}(\alpha,\beta)$ with $\alpha \leq \left\lceil \lg \left(2(4n^2+1)\right) \right\rceil$ and $\beta \leq 2^{\alpha}m+2 = O(n^3)$. It follows that each gadget $\operatorname{star}(\alpha,\beta)$ has at most $O(n^3 \log n)$ vertices. Because we use at most 2n gadgets, H contains a total of at most $n + 2nO(n^3 \log n) = O(n^4 \log n)$ vertices.

5. Complexity of Pebbling Number

Although the optimal pebbling number has received some study, combinatorialists have focused more attention on the pebbling number. Recall that the *r*-pebbling number $\pi(G, r)$ is the minimum k such that r is reachable under every distribution of size k. Similarly, the pebbling number $\pi(G)$ is the minimum k such that every vertex is reachable under every distribution of size k. It is clear from the definitions that if n is the number of vertices in G, then $\hat{\pi}(G) \leq n \leq \pi(G)$. At first glance, it may not be clear that $\pi(G)$ is well defined. In fact, if G is not connected, then we can place arbitrarily many pebbles in a single component and we will not be able to place pebbles on vertices outside the component. However, for connected graphs, $\pi(G)$ is well defined; we implicitly assume that G is connected. Indeed, if d is the diameter of G, every vertex is reachable provided that our distribution is forced to place at least 2^d pebbles on some vertex. We record this observation as a proposition.

Proposition 5.1. Let G be a graph with diameter d. Then $\pi(G) \leq (2^d - 1)n + 1$.

We call the problem of deciding whether $\pi(G, r) \leq k$ R-PEBBLING-NUMBER (abbreviated RPN); similarly, we define PEBBLING-NUMBER (abbreviated PN) to be the problem of deciding whether $\pi(G) \leq k$. In this section, we establish that PN and RPN are Π_2^{P} -complete. First, note that both languages are in Π_2^{P} . Indeed, to decide if $\pi(G) \leq k$, our machine need only check that for all distributions p of size k and all target vertices r, there exists an orderable digraph $D_{p,r}$ that places a pebble on r. The distributions of size k, the target vertices, and the digraphs $D_{p,r}$ are all describable using poly $(n, \log k)$ bits. Further, ORDERABLE is in P. It follows that PN is in Π_2^{P} . A similar argument shows that RPN is in Π_2^{P} . The seminal Π_2^{P} -complete problem is a quantified version of 3SAT whose instances consist of a 3CNF formula ϕ over a set of universally quantified variables and a set of existentially quantified variables (see [P]). We say that ϕ is *valid* if for every setting of the universally quantified variables, there is a setting of the existentially quantified variables which satisfies ϕ . The decision problem $\forall \exists 3SAT$ is to determine whether ϕ is valid.

Just as 3SAT remains NP-complete when ϕ is restricted to be in canonical form (recall Definition 3.1), $\forall \exists 3$ SAT remains Π_2^{P} -complete when ϕ is restricted to be in canonical form. We call this restriction $\mathbb{R}\forall \exists 3$ SAT.

We show that RPN is $\Pi_2^{\rm P}$ -complete by a reduction from $\mathbb{R}\forall\exists 3$ SAT. Whereas our reduction to OPN produces graphs H with the property that only one distribution can possibly succeed in witnessing $\hat{\pi}(H) \leq k$, our reduction to RPN produces graphs with the property that almost all distributions succeed in being able to place a pebble on r. It is the rare "difficult" distributions – those which may not allow a pebble to be placed on r – that correspond to settings of the universally quantified variables in our $\mathbb{R}\forall\exists 3$ SAT formula. Given a distribution of k pebbles to the graph we produce, either r is easily reachable, or the distribution corresponds to a setting f of the universally quantified variables in ϕ and r is reachable if and only if ϕ is satisfiable under f.

Our reduction from $R \forall \exists 3SAT$ to RPN involves the construction of several graphs, each building on the previous construction. We refer to the *i*th graph we produce as $G_i = G_i(\phi)$. We present the reduction with respect to a fixed instance ϕ of $R \forall \exists 3SAT$.

5.1. The Underlying Graph. We obtain G_1 from ϕ by modifying $G^{\text{NPR}}(\phi)$ slightly. That is, for each universally quantified variable x_i in ϕ , we remove both edges from the variable gadget X_i in $G^{\text{NPR}}(\phi)$ associated with x_i and remove one pebble each from the endpoints of X_i , so that the endpoints of X_i start with one pebble instead of two. (We leave intact variable gadgets X_j corresponding to existentially quantified variables x_j in ϕ .) Let $n_1 = n(G_1)$ be the number of vertices in G_1 , let $e_1 = e(G_1)$, and let p_1 be the distribution on G_1 . The following definition gives the correspondence between settings of the universally quantified variables in ϕ and distributions of pebbles in G_1 .

Definition 5.2. For each setting f of the universally quantified variables in ϕ , let $p_{1,f}$ be the distribution of pebbles to G_1 given by adding the following pebbles to p_1 . For each x_i with $f(x_i) =$ true, add one pebble to each of the two vertices associated with positive occurrences of x_i in ϕ . For each x_i with $f(x_i) =$ false, add two pebbles to the vertex associated with the negative instance of x_i .

Observe that under any $p_{1,f}$, each vertex in G_1 contains at most two pebbles. Our interest in G_1 under the distributions $p_{1,f}$ is based on the following proposition, whose proof is similar to that of Theorem 3.3.

Proposition 5.3. There is a nonrepetitive sequence of pebbling moves which places a pebble on r in G_1 starting from $p_{1,f}$ if and only if there is a setting of the existentially quantified variables in ϕ which, together with f, satisfies ϕ .

Let t be the number of pebbles in the $p_{1,f}$. Because $p_{1,f}$ assigns at most two pebbles to each vertex in $G_1, t \leq 2n_1$. We obtain G_2 from G_1 by setting $\alpha = \max \{ \lg 2t, 4 \lg e_1 \}$ and replacing each edge in G_1 with a path of length $\alpha + 1$; that is, $G_2 = \mathcal{S}(G_1, \alpha)$ (recall Definition 3.6). Let n_2 be the number of vertices in G_2 .

Let p_2 be the distribution of pebbles to G_2 so that p_2 and p_1 agree on all vertices in G_1 and $p_2(v) = 1$ for all vertices v introduced in our construction of G_2 from G_1 . Similarly, let $p_{2,f}$ be the distribution of pebbles to G_2 so that $p_{2,f}$ and $p_{1,f}$ agree on all vertices in G_1 and $p_{2,f}(v) = 1$ for all vertices v introduced in our construction of G_2 from G_1 .

We call G_2 the underlying graph and a distribution $p_{2,f}$ an underlying distribution. Observe that by Lemma 3.7, there is a nonrepetitive sequence of pebbling moves which places a pebble on r in G_1 under $p_{1,f}$ if and only if there is an arbitrary sequence of pebbling moves in G_2 under $p_{2,f}$ which places a pebble on r. Together with Proposition 5.3, we obtain the following.

Proposition 5.4. There is a sequence of pebbling moves which places a pebble on r in G_2 starting from $p_{2,f}$ if and only if there is a setting of the existentially quantified variables in ϕ , which, together with f, satisfies ϕ .

One useful property of the underlying graph together with an underlying distribution is that it is not possible to accumulate more than five pebbles on any vertex. This property will be instrumental in arguing that the gadgets we attach to the underlying graph behave correctly.

Proposition 5.5. It is not possible to place more than five pebbles on any vertex in G_2 starting from any $p_{2,f}$.

Proof. Suppose for a contradiction it is possible to place at least six pebbles on a vertex u in G_2 . First, suppose u is a vertex introduced in our construction of G_2 from G_1 , so that u is an internal vertex w_i , $1 \leq i \leq \alpha$, in a one use path $P = w_0 w_1 \dots w_\alpha w_{\alpha+1}$. Let D be a signature of a minimum sequence of pebbling moves which places at least six pebbles on v_i . Because $p_{2,f}(w_i) = 1$, for balance $(D, p_{2,f}, w_i) \geq 6$ we must have that the indegree of w_i is at least five. It follows by the pigeonhole principle that either the multiplicity of $w_{i-1}w_i$ or $w_{i+1}w_i$ is at least three. If the former is true, we can apply Corollary 3.5 to obtain a sequence of pebbling moves that places $2^i(3-1) + 2 \cdot 1 \geq 6$ pebbles on w_0 . Similarly, if the latter is true, we apply Corollary 3.5 to obtain a sequence of pebbling moves that places $2^{\alpha-i}(3-1) + 2 \cdot 1 \geq 6$ pebbles on $w_{\alpha+1}$. Because w_0 and $w_{\alpha+1}$ are vertices in G_1 , it suffices to show that it is not possible to place more than five pebbles on any vertex in G_1 .

Suppose that u is in G_1 . Because it is possible to place at least six pebbles on u in G_2 starting from $p_{2,f}$, by Lemma 3.7, there is a nonrepetitive sequence of pebbling moves that places at least six pebbles on u in G_1 starting from $p_{1,f}$. But this is clearly impossible, because the maximum degree in G_1 is three and each vertex receives at most two pebbles from $p_{1,f}$.

Now that we have established the important properties of the underlying graph and the underlying distributions, we attach gadgets to the vertices in the underlying graph. Just as the star gadgets we attach in our reduction from DPR to OPN force any potentially successful distribution to take a certain form, our gadgets here force any potentially unsuccessful distribution to take a form which effectively induces one of the underlying distributions on the underlying graph.

5.2. The Gadgets. We introduce three classes of gadgets: the null gadget, the fork gadget, and the eye gadget. In this section, we explore the relevant properties of our gadgets as isolated graphs.

All classes of gadgets share some common properties. The gadgets have *attachment vertices*; later, we will attach gadgets to the underlying graph by identifying the attachment vertices of a gadget with vertices in the underlying graph. A *supply quota s* assigns each attachment vertex v a number s(v); each gadget has one or more supply quotas. Under a particular distribution q, a gadget *satisfies s* if q covers s.

The gadgets have overflow vertices, which are adjacent to r; we call the edges between the overflow vertices and r the overflow edges. We say that a gadget has an overflow threshold of k if r is reachable via an overflow edge under every distribution of size k.

Let q be a distribution of pebbles to a gadget. If the gadget is able to satisfy any one of its supply quotas, or if r is reachable via an overflow edge, we say that the gadget is *potent* under q. We say that a gadget has a *potency threshold* of k if the gadget is potent under every distribution of k pebbles.

Every gadget has one or more *critical distributions*, each of equal size. If q is a critical distribution and s is a supply quota, we say that q breaches s if there exists a vertex v such that it is possible to place more than s(v) pebbles on v starting from q.

Our critical distributions and supply quotas are in bijective correspondence; that is, for each critical distribution there is a corresponding supply quota and vice versa. Each critical distribution q exhibits the following critical distribution properties:

- (1) starting from q, r is not reachable via an overflow edge
- (2) q does not breach its corresponding supply quota

As we present the gadgets, their supply quotas, and their critical distributions, we will establish an overflow threshold, a potency threshold, and the critical distribution properties.

To motivate the study of these parameters, we outline their use in our proof of the correctness of our reduction. Given an $R \forall \exists 3SAT$ instance ϕ , we compute H and k such that ϕ is valid if and only if $\pi(H, r) \leq k$. We construct H by attaching various gadgets to the underlying graph and we set k to be the sum, over all gadgets, of the size of the gadget's critical distributions.

Suppose that ϕ is valid and consider a distribution of k pebbles to H. If some gadget is assigned fewer pebbles than its potency threshold, the pigeonhole principle implies that some gadget receives more pebbles than its overflow threshold, and hence r is reachable. Otherwise, all gadgets are potent. If r is reachable via some gadget's overflow edge, we are done. Otherwise, every gadget is able to satisfy one of its supply quotas;



FIGURE 5.1. The null gadget. The dashed line represents a path of length c, the circle around v indicates that v is an attachment vertex, and the box around w indicates that w is an overflow vertex.

this implies that our initial distribution on H covers some $p_{2,f}$. Because ϕ is satisfiable under f, we obtain from Proposition 5.4 a sequence of pebbling moves in the underling graph which places a pebble on r.

The converse direction is somewhat trickier, but proceeds roughly as follows. Suppose that $\pi(H,r) \leq k$ and consider a setting f of the universally quantified variables of ϕ . We assign pebbles to H by selecting (according to f) a critical distribution for each gadget. Because $\pi(H,r) \leq k$, we obtain a signature Dof a minimum sequence of pebbling moves which places a pebble on r. Next, we argue that our critical distribution properties still apply even though the gadgets have been attached to the underlying graph. Then we show how D can be used to obtain a sequence of pebbling moves in the underlying graph starting from $p_{2,f}$ which places a pebble on r. A final application of Proposition 5.4 implies that ϕ is satisfiable under f.

Our gadgets are defined in terms of two parameters, β and c. We set c = 3 (in fact, any constant c so that 2^c exceeds the constant obtained in Proposition 5.5 will do). We postpone fixing the precise value of β ; suffice it to say we will choose $\beta = \Theta(\log n_2)$. Our gadgets use small paths of length c to provide some separation between the underlying graph and more sensitive areas of our gadgets. We use larger paths of length β so that the number of pebbles in a gadget's critical distribution far exceeds its potency threshold.

5.2.1. The Null Gadget. The null gadget is a path of length c and appears in Figure 5.1 on page 16. We use the null gadget to ensures that every vertex in the underlying graph is not too far away from r, so that distributions which concentrate pebbles on the underlying graph quickly imply that r is reachable. The null gadget has a single supply quota s, with s(v) = 0; its corresponding critical distribution q assigns zero pebbles to each vertex in the null gadget.

- **Overflow threshold:** Because c is a fixed constant, the null gadget is a fixed graph which does not depend upon ϕ . By Proposition 5.1, its pebbling number is a fixed constant, say a, not depending upon ϕ . Clearly, if there are 2a pebbles in the null gadget, then it is possible to place two pebbles on w and hence one pebble on r. It follows that 2a = O(1) is an overflow threshold for the null gadget. **Potency threshold:** Because s is trivially satisfied, the null gadget has a potency threshold of 0.
- Critical distribution properties: Because q assigns zero pebbles to the null gadget, it is clear that under q, the null gadget does not breach s, nor is it possible to place a pebble on r via the null gadget's overflow edge.

5.2.2. The Fork Gadget. The fork gadget consists of three paths P_1, P_2, P_3 which share only a common endpoint, as shown in Figure 5.2 on page 17. The fork gadget is responsible for injecting one pebble in the underlying graph at the attachment location, much like the star gadgets in the previous section. It has one supply quota s with s(v) = 1; the corresponding critical distribution q is given by $q(u) = 2 \cdot 2^{\beta+c} - 1$ and q(x) = 0 for all $x \neq u$.

- **Overflow threshold:** The fork gadget has an overflow threshold of $2 \cdot 2^{\beta+c} + O(1)$. Indeed, if the fork gadget is unable to place two pebbles on w (and hence one on r), there can be at most O(1) pebbles on P_2 and P_3 . Secondly, there can be at most $2 \cdot 2^{\beta+c} 1$ pebbles in P_1 and P_2 . It follows that the fork gadget can contain at most $2 \cdot 2^{\beta+c} + O(1)$ pebbles if r is not reachable via an overflow edge.
- **Potency threshold:** The fork gadget has a potency threshold of $2^{\beta+c} + O(1)$. Indeed, if the fork gadget is not potent, then it must have at most O(1) pebbles on P_2 , or else it would be able to place a pebble on r. Similarly, it must have at most $2^{\beta+c} 1$ pebbles on P_1 and P_3 , or else it would be able to place a pebble on v_3 and therefore satisfy s.
- Critical distribution properties: Both the standard weight equation and the greedy tree lemma show that under q, the fork gadget does not breach s, nor is r reachable via an overflow edge.



FIGURE 5.2. The fork gadget. The dashed lines represent paths P_2, P_3 of length c, the solid line represents a path P_1 of length β , the circle around v indicates that v is an attachment vertex, and the box around w indicates that w is an overflow vertex.



FIGURE 5.3. The eye gadget. The dashed lines represent paths of length c, and the solid lines represent paths P_0, P_1, P_2, P_3 of length β . The circled vertices v_i are attachment vertices; the boxed vertices w_i are overflow vertices.

5.2.3. The Eye Gadget. The eye gadget is the most complex of our three gadgets, and it is at the heart of our reduction. Our reduction attaches one eye gadget for each universally quantified variable in ϕ . The eye gadget is shown in Figure 5.3 on page 17.

The eye gadget has two supply quota/critical distribution pairs. The pair (s^+, q^+) corresponds to a positive (true) setting of the variable x and the pair (s^-, q^-) corresponds to a negative (false) setting of x. We call s^+ the positive supply quota and we call s^- the negative supply quota. Similarly, we call q^+ the positive critical distribution and q^- the negative critical distribution.

We define the supply quotas via $s^+(v_1) = s^+(v_3) = 1$, $s^+(v_2) = 0$, and $s^-(v_1) = s^-(v_3) = 0$, $s^-(v_2) = 2$. Similarly, the critical distributions are given by $q^+(u_1) = q^+(u_3) = 2 \cdot 2^{\beta+c} - 1$, $q^+(u_0) = q^+(u_2) = 2^{\beta+c} - 1$, and $q^-(u_1) = q^-(u_3) = 2^{\beta+c} - 1$, $q^-(u_0) = q^-(u_2) = 2 \cdot 2^{\beta+c} - 1$.

Let F be the subgraph of the eye gadget obtained by removing the u_i and all interior vertices of paths the P_i . Observe that F depends only on c and therefore, like the null gadget, F is a fixed graph, not depending upon ϕ . It follows that $\pi(F) = O(1)$.

- **Overflow threshold:** The eye gadget has an overflow threshold of $6 \cdot 2^{\beta+c} + O(1)$. Suppose the eye gadget contains k pebbles and it is not possible to place a pebble on r via one of the overflow edges. We show that $k \leq 6 \cdot 2^{\beta+c} + O(1)$. Immediately, we have that F contains at most $2\pi(F) = O(1)$ pebbles, or else it would be possible to place two pebbles on w_0 and hence one pebble on r. To bound the number of pebbles in the P_i , we consider two cases. First, suppose that each P_i contains fewer than $2^{\beta+c}$ pebbles; in this case, we have that $k \leq 4 \cdot 2^{\beta+c} + O(1)$. Otherwise, suppose that P_j has at least $2^{\beta+c}$ pebbles. Clearly, P_j has at most $2 \cdot 2^{\beta+c} 1$ pebbles, or else we could use these pebbles to place a pebble on r via the overflow vertex w_j ; similarly, the opposite path P_{j+2} contains at most $2 \cdot 2^{\beta+c} 1$ pebbles (subscript arithmetic is understood modulo 4). Finally, the remaining paths P_{j-1}, P_{j+1} each contain at most $2^{\beta+c} 1$ pebbles; indeed, if P_{j-1} (P_{j+1}) contained $2^{\beta+c}$ pebbles, we could use them to place one pebble on w_{i-1} (w_i) and we could use $2^{\beta+c} 4$ pebbles, and so $k \leq 6 \cdot 2^{\beta+c} + O(1)$.
- Potency threshold: The eye gadget has a potency threshold of $5 \cdot 2^{\beta+c} + O(1)$. Suppose the eye gadget contains k pebbles, r is not reachable via an overflow edge, and it is not possible to satisfy s^+ or s^- . We show that $k \leq 5 \cdot 2^{\beta+c} + O(1)$. As before, we have that F contains at most O(1) pebbles. To bound the number of pebbles in the P_i , we consider the same two cases as before. If each path has fewer than $2^{\beta+c}$ pebbles, we immediately have $k \leq 4 \cdot 2^{\beta+c} + O(1)$ and we're done. Otherwise, suppose P_j has at least $2^{\beta+c}$ pebbles. Once again, we have that P_j contains at most $2 \cdot 2^{\beta+c} 1$ pebbles, and P_{j-1}, P_{j+1} each contain at most $2^{\beta+c} 1$. However, now the opposite path P_{j+2} has at most $2^{\beta+c} 1$ pebbles. Indeed, if P_j, P_{j+2} both contain at least $2^{\beta+c}$ pebbles, then we can either place one pebble each on v_1 and v_3 , satisfying s^+ (as is the case if $\{j, j+2\} = \{1,3\}$), or we can place two pebbles on v_2 , satisfying s^- (as is the case if $\{j, j+2\} = \{0,2\}$). It follows that the paths contain at most $5 \cdot 2^{\beta+c} 4$ pebbles, implying $k \leq 5 \cdot 2^{\beta+c} + O(1)$.
- **Critical distribution properties:** It remains to verify the critical distribution properties for q^+ and q^- . First, we show that under $q \in \{q^+, q^-\}$, r is not reachable via an overflow vertex. Let R be the set of overflow vertices in the eye gadget, and let D be the signature of a minimum sequence of pebbling moves that places two pebbles on a vertex in R. By Lemma 2.8, we have that each vertex in R has outdegree zero in D. Observe that deleting R from the eye gadget results in a graph with three components; let A_1 be the component containing P_1 , let A_2 be the component containing P_0 and P_2 , and let A_3 be the component containing P_3 . Let $T_1 = A_1 + \{w_0, w_1\}$, let $T_2 = A_2 + R$, and let $T_3 = A_3 + \{w_2, w_3\}$, as shown in Figure 5.4 on page 19. Let D_l be the digraph obtained by deleting from D all pebbling moves outside of T_l . Because D is acyclic, it is immediate that each D_l is acyclic. Observe that for all $x \in V(T_l) R$, we have balance $(D_l, q, x) =$ balance (D, q, x). Furthermore, because $d_D^+(w_i) = 0$ for all $w_i \in R$, we have that $d_{D_l}^+(w_i) = 0$. Therefore by the acyclic orderability characterization, D_l is orderable.

Let w_i be the overflow vertex on which D places two pebbles. Because $q(w_i) = 0$, we have that the indegree of w_i in D is at least two. Suppose two edges into w_i are contained in the same tree T_l . Then D_l is the signature of a sequence of pebbling moves in T_l starting from q that places at least two pebbles on w_i . By the greedy tree lemma, the greedy pebbling strategy in T_l under q places at least two pebbles on w_i . However it is easily checked that regardless of $q \in \{q^+, q^-\}, T_l \in \{T_1, T_2, T_3\}$, and $w_i \in R$, the greedy strategy in T_l under q places at most one pebble on w_i . Alternatively, suppose that D contains edges into w_i from two distinct trees. Because w_i is in T_2 and one other tree, it must be that D contains an edge into w_i from T_2 . Then D_2 is a signature of a sequence of pebbling moves in T_2 starting from q which places a pebble on w_i ; therefore the greedy strategy in T_2 starting from q^+ is unable to place any pebbles on any overflow vertex, it follows that $q = q^-$. Suppose that D contains an edge into w_i ; therefore the greed of pebbling moves in T_l starting from q^- places a pebble on w_i ; therefore the greedy strategy in T_l starting from q^- places a pebble on w_i ; therefore the greedy strategy in T_l starting from q^- places a pebble on w_i ; therefore the greedy strategy in T_l starting from q^- places a pebble on w_i ; therefore the greedy strategy in T_l starting from q^- places a pebble on w_i . But now a familiar contradiction is at hand: it is easily checked that regardless of $T_l \in \{T_1, T_3\}$ and $w_i \in R$, the greedy strategy in T_l starting from q^- is unable to place a pebble on w_i .

Let $(s,q) \in \{(s^+,q^+), (s^-,q^-)\}$. It remains to show that q does not breach s. Suppose for a contradiction that D is the signature of a minimum sequence of pebbling moves which witnesses



FIGURE 5.4. The overflow vertices split the eye gadget into three trees T_1, T_2, T_3

that q breaches s. We have that the outdegree of each overflow vertex $w_i \in R$ is zero; indeed, if $d_D^+(w_i) \geq 1$, then by Lemma 2.6 we would obtain a sequence of pebbling moves placing two pebbles on w_i , a contradiction. As before, let D_l be the digraph obtained from D by deleting all edges outside of T_l ; as before, we have that D_l is orderable in T_l . It follows that if D places more than $s(v_l)$ pebbles on v_l , then D_l witnesses that it is possible to place more than $s(v_l)$ pebbles on v_l in T_l starting from q. By the greedy tree lemma, the greedy strategy places more than $s(v_l)$ pebbles of $(s,q) \in \{(s^+,q^+), (s^-,q^-)\}$ and $l \in \{1,2,3\}$, the greedy strategy in T_l starting from q places exactly $s(v_l)$ pebbles on v_l .

5.2.4. Summary. We summarize the various parameters of our gadgets in the following table.

gadget	potency threshold	size of critical distributions	overflow threshold
null	0	0	O(1)
fork	$2^{\beta+c} + O(1)$	$2 \cdot 2^{\beta+c} - 1$	$2 \cdot 2^{\beta+c} + O(1)$
eye	$5 \cdot 2^{\beta+c} + O(1)$	$6 \cdot 2^{\beta+c} - 4$	$6 \cdot 2^{\beta+c} + O(1)$

From the table, we obtain the gap lemma.

Lemma 5.6 (Gap Lemma). There exists a nonnegative constant C (depending only on c) such that for each gadget, the overflow threshold exceeds the size of the critical distributions by at most C, and for the fork and eye gadgets, the size of the critical distributions exceed the potency threshold by at least $2^{\beta+c} - C$.

5.3. Construction of *H*. We set $\beta = \lceil \lg 3Cn_2 \rceil$, with *C* as in Lemma 5.6.

Armed with our gadgets and our underlying graph G_2 , we are able to describe the last step in our reduction from R $\forall \exists 3$ SAT to RPN. For each pebble in p_2 on a vertex z in the underlying graph, we attach a fork gadget to z by identifying the attachment vertex v in the fork gadget with z. For each triplet z_1, z_2, z_3 of vertices in G_2 corresponding to a universally quantified variable x in ϕ , with z_1, z_3 corresponding to positive occurrences of x in ϕ and z_2 corresponding to the negative occurrence of x in ϕ , we attach an eye gadget by identifying the attachment vertex v_i in the eye gadget with z_i in the underlying graph. Finally, for any vertex $z \neq r$ in the underlying graph to which we did not attach a fork or eye gadget, we attach a null gadget by identifying v in the null gadget with z in the underlying graph. Let H be the resulting graph, and let k be the sum, over all gadgets in H, of the size of the gadget's critical distributions. Our reduction from $R \forall \exists 3SAT$ to RPN outputs H, k, and r.

Note that we attach gadgets to the underlying graph by identifying attachment vertices in gadgets with vertices in the underlying graph, so that in H, each attachment vertex v is a member of the underlying graph and also a member of a gadget. Furthermore, by our construction, every vertex other than r in the underlying graph is identified with an attachment vertex, so the vertices in the underlying graph are exactly the attachment vertices together with r.

We pause to observe two important properties about H.

Proposition 5.7. In constructing H, we attach at most two gadgets to every vertex in the underlying graph.

Proof. Recall that p_2 assigns at most two pebbles to any vertex in the underlying graph; furthermore, p_2 assigns at most one pebble to any vertex associated with a universally quantified variable in ϕ .

Proposition 5.8. The diameter of H is at most $2\beta + O(1)$.

Proof. It suffices to show that for each z in H, the distance from z to r is at most $\beta + O(1)$. If $z \neq r$, then z is contained in some gadget. In each gadget, every vertex is at most $\beta + O(1)$ from an overflow vertex. \Box

5.4. R-PEBBLING-NUMBER is Π_2^{P} -complete.

Proposition 5.9. Let $a_1, \ldots, a_n, b_1, \ldots, b_n$, and x be real numbers with $\sum_{i=1}^n a_i \ge \sum_{i=1}^n b_i$. If $a_n < b_n - x$ then there exists i such that $a_i > b_i + x/(n-1)$.

Proof. By contradiction. Otherwise,

$$\sum_{i=1}^{n} a_{i} = \sum_{i=1}^{n-1} a_{i} + a_{n}$$

$$< \left(\sum_{i=1}^{n-1} b_{i} + \frac{x}{n-1}\right) + b_{n} - x$$

$$< \sum_{i=1}^{n} b_{i}.$$

We have accumulated the tools needed to show the correctness of our reduction.

Theorem 5.10. ϕ is valid if and only if $\pi(H, r) \leq k$.

Proof. (\implies). Suppose that ϕ is valid and let p be a pebble distribution on H of size k. We may assume p(r) = 0. Let l be the number of gadgets in H, label the gadgets as Q_1, \ldots, Q_l , let a_i be the number of pebbles that p assigns to Q_i , and let b_i be the size of Q_i 's critical distributions. Because every vertex in H besides r belongs to at least one gadget, we have $\sum_{i=1}^{n} a_i \ge k = \sum_{i=1}^{n} b_i$.

We consider several cases. First, suppose there is some gadget Q_i to which p assigns fewer pebbles than Q_i 's potency threshold; by the gap lemma, we have that $a_i < b_i - (2^{\beta+c} - C)$. By Proposition 5.9, there is some Q_j to which p assigns at least $(2^{\beta+c} - C)/(l-1)$ pebbles more than Q_j 's overflow threshold. By Proposition 5.7, $l-1 \leq l \leq 2n_2$. It follows that Q_j contains at least

$$\begin{array}{rcl} \frac{2^{\beta+c}-C}{2n_2} & \geq & \frac{2^{\beta}-C}{2n_2} \\ & \geq & \frac{3Cn_2-C}{2n_2} \\ & \geq & \frac{2Cn_2}{2n_2} \\ & \geq & C \\ & & & \\ & & & \\ & & & 20 \end{array}$$

more pebbles than the size of its critical distributions. It follows from Lemma 5.6 that Q_j contains at least as many pebbles as its overflow threshold and therefore we can place a pebble on r via one of Q_j 's overflow edges. Otherwise, p assigns every gadget at least as many pebbles as its potency threshold. If there is some gadget which is able to place a pebble on r via an overflow edge, then we are done. Otherwise, for every gadget Q, there is a supply quota s such that Q under p satisfies s. Using these supply quotas, we obtain a setting f of the universally quantified variables in ϕ as follows. We set f(x) = true if the eye gadget associated with x satisfies its positive supply quota s^+ ; otherwise, the eye gadget associated with x must meet the negative supply quota s^- and we set f(x) = false. We claim that p covers $p_{2,f}$. In each gadget, execute the pebbling moves witnessing that the gadget satisfies its supply quota. The fork gadgets alone produce a distribution that is at least as good as p_2 , and the eye gadgets supply the additional pebbles proscribed by $p_{2,f}$. Because ϕ is valid, it follows from Proposition 5.4 that r is reachable.

(\Leftarrow). Suppose that $\pi(H, r) \leq k$ and let f be a setting of the universally quantified variables in ϕ . We obtain a setting of the existentially quantified variables in ϕ witnessing that ϕ is satisfiable under f. Naturally, we study a pebble distribution p on H of size k corresponding to f; we construct p by choosing a critical distribution q_i for each gadget Q_i . If Q_i is not an eye gadget, then Q_i has only one critical distribution and our selection of q_i is forced. If Q_i is an eye gadget, we let q_i be the positive critical distribution q^+ if f(x) = true and we let q_i be the negative critical distribution q^- otherwise. Note that p does not assign any pebbles to any vertex in the underlying graph. Let s_i be the supply quota associated with q_i .

Let H' be the graph obtained from H by removing all the overflow edges. Our first task is to establish the analog of Proposition 5.5 for H'.

Claim 5.11. In H' starting from p, it is not possible to place more than five pebbles on any vertex in the underlying graph.

Proof. Suppose for a contradiction that D is the signature of a minimum sequence of pebbling moves that places at least six pebbles on some vertex w in the underlying graph.

Claim. D does not contain an edge whose origin is inside the underlying graph and whose destination is outside the underlying graph.

Proof. Suppose for a contradiction that uv is an edge in D from a vertex u in the underlying graph to some vertex v not in the underlying graph. Because H' does not contain any overflow edges, it must be that uv is an edge on a path of length c in some gadget; let this path be $P = x_0 \dots x_c$, with $u = x_c$ and $v = x_{c-1}$. It follows that D contains the edge x_1x_0 , or else D contains a cycle or a proper sink other than w, contradicting the minimum signatures lemma. Because $p(x_i) = 0$ for each internal vertex of P, we have by Corollary 3.5 that it is possible to place $2^c = 8 \ge 6$ pebbles on u using fewer pebbling moves, a contradiction. Therefore D does not contain an edge from the underlying graph to a vertex outside the underlying graph.

Claim. For each u in the underlying graph, the number of edges in D into u with origins outside the underlying graph is at most $p_{2,f}(u)$.

Proof. If this were not the case, then there is some gadget Q_i attached to u such that D contains more than $s_i(u)$ edges from Q_i into u. Construct D' from D by deleting all edges not contained in Q_i . Clearly, $D' \subseteq D$ is acyclic; we show that D' is orderable by verifying the balance condition. Consider a vertex v in Q_i . Recall that H' does not contain overflow edges, and therefore if v is not an attachment vertex, then the neighborhood of v is contained in Q_i . It follows that if v is not an attachment vertex, we have balance $(D', q_i, v) =$ balance (D, p, v). Alternatively, if v is an attachment vertex, we have that $d_{D'}^+(v) = 0$, or else D' (and hence D) would contain an edge from a vertex v in the underlying graph to a vertex outside the underlying graph, contradicting our previous claim. It follows that if v is an attachment vertex, we have that $d_{D'}^+(v) \ge 0$. By the acyclic orderability characterization, we have that D' is orderable under q_i . Together with $d_{D'}^-(u) > s_i(u)$ and $d_{D'}^+(u) = 0$ (recall u is an attachment vertex), we have that balance $(D', q_i, u) > s_i(u)$. Therefore D' witnesses that it is possible to place more than $s_i(u)$ pebbles on u in Q_i starting from q_i , contradicting Q_i 's critical distribution properties.

We return to our proof of Claim 5.11. Construct D' from D by removing all edges from D that are not in the underlying graph. Clearly, $D' \subseteq D$ is acyclic. We show that D' is orderable under $p_{2,f}$ by checking the balance condition. For each u in the underlying graph, we have balance $(D', p_{2,f}, u) \geq$ balance(D, p, u). Indeed, at most $p_{2,f}(u)$ edges into u are deleted from D in our construction of D'; however, p(u) = 0, so that $p_{2,f}$ offsets this decrease in balance. It follows that D' is orderable under $p_{2,f}$. Together with balance $(D', p_{2,f}, w) \ge$ balance $(D, p, w) \ge 6$, we have that it is possible to place at least six pebbles on w starting from $p_{2,f}$ in the underlying graph, contradicting Proposition 5.5. This completes our proof of Claim 5.11.

We return to our proof of Theorem 5.10. Let D be the signature of a minimal sequence of pebbling moves in H starting from p that places a pebble on r.

Claim 5.12 (No Backflow into Gadgets Claim). D does not contain an edge from a vertex inside the underlying graph to a vertex outside the underlying graph.

Proof. By the minimum signatures lemma, we have that D contains at most one pebbling move along an overflow edge and any such pebbling move must be directed from an overflow vertex into r. Construct D' from D by removing this edge if it exists. Because r has outdegree zero in D, the acyclic orderability characterization implies that D' is orderable. Furthermore, because D' does not contain any pebbling move along overflow edges, D' yields a sequence of pebbling moves in H'.

Because D' is constructed from D by removing at most one edge into r, it suffices to show that D' does not contain an edge from a vertex inside the underlying graph to a vertex outside the underlying graph. Suppose for a contradiction that D' contains an edge uv from u inside the underlying graph to v outside the underlying graph. It must be that uv is a pebbling move along a path P of length c in some gadget. Let $P = x_0 \dots x_c$ with $u = x_c$ and $v = x_{c-1}$. It follows that D' contains the edge x_1x_0 . Indeed, if D' does not have x_1x_0 as an edge, neither does D (after all, $x_0 \neq r$), and so D contains a cycle or a proper sink other than r, contradicting the minimum signatures lemma. Therefore D' contains the pebbling move x_1x_0 .

Recalling that p assigns each internal vertex of P zero pebbles, Lemma 2.6 implies that there is an orderable $D'' \subseteq D'$ which places at least $2^c = 8$ pebbles on $x_c = u$. But now D'' is a signature witnessing that it is possible to place at least six pebbles on u in H' starting from p, contradicting Claim 5.11.

Let us resume our proof of Theorem 5.10. Construct D_i from D by deleting from D all edges not contained in Q_i or along Q_i 's overflow edges.

Claim 5.13. D_i is orderable under q_i , and for each attachment vertex v, balance $(D_i, q_i, v) = d_{D_i}^-(v)$.

Proof. Because $D_i \subseteq D$, D_i is acyclic and so it suffices to verify the balance condition. Because $d_D^+(r) = 0$, clearly $d_{D_i}^+(r) = 0$ and so the balance condition is satisfied at r. Consider a vertex v in Q_i . Unless v is an attachment vertex, all edges incident to v in D also appear in D_i , and so balance $(D_i, q_i, v) =$ balance(D, p, v). Otherwise, if v is an attachment vertex, then $d_{D_i}^+(v) = 0$ or else D would contain an edge from a vertex in the underlying graph to a vertex outside the underlying graph, contradicting Claim 5.12. Together with $q_i(v) = 0$, it follows that balance $(D_i, q_i, v) = d_{D_i}^-(v)$. By the acyclic orderability characterization, D_i is orderable under q_i .

Claim 5.14. For each u in the underlying graph, D contains at most $p_{2,f}(u)$ edges from outside the underlying graph into u.

Proof. Suppose that u is a counterexample to the claim. If u = r, then there is some gadget Q_i such that D contains an edge wr into r along one of Q_i 's overflow edges. But D_i also contains wr and, by Claim 5.13, D_i is orderable under q_i . Clearly, balance $(D_i, q_i, r) \ge 1$ and therefore r is reachable in Q_i under q_i , contradicting the critical distribution properties of Q_i . Otherwise, if $u \ne r$, then there is some gadget Q_i such that D contains more than $s_i(u)$ edges into u from vertices in Q_i . But these edges are also in D_i , so that $d_{D_i}(u) > s_i(u)$. By Claim 5.13, D_i is the signature of a sequence of pebbling moves in Q_i under q_i placing more than $s_i(u)$ pebbles on u, contradicting Q_i 's critical distribution properties.

Let us complete our proof of Theorem 5.10. Construct E from D by deleting from D any edges outside the underlying graph. We show that E is orderable under $p_{2,f}$. Clearly, $E \subseteq D$ is acyclic and therefore it suffices to check the balance condition. Consider a vertex u in the underlying graph, and let m be the number of edges into u from outside the underlying graph. In constructing E from D, the balance of udecreases by m; by Claim 5.14, we have $m \leq p_{2,f}(u)$. Because p(u) = 0, changing distributions from p to $p_{2,f}$ increases the balance of u by $p_{2,f}(u)$. It follows that $\text{balance}(E, p_{2,f}, u) \geq \text{balance}(D, p, u)$. Therefore E is orderable under $p_{2,f}$ and so r is reachable in the underlying graph under $p_{2,f}$. A final application of Proposition 5.4 implies that ϕ is satisfiable under f. This completes our proof of Theorem 5.10.

We are now able to complete our proof that R-PEBBLING-NUMBER is $\Pi_2^{\rm P}$ -complete.

Theorem 5.15. R-PEBBLING-NUMBER is Π_2^{P} -complete, even when the diameter of H is at most $O(\log n(H))$ and k = poly(n(H)).

Proof. We have already observed that RPN is in Π_2^P and checked the correctness of our reduction; it remains to check the diameter condition on H and that H and k are not too large relative to ϕ so that our reduction is computable in polynomial time. By Proposition 5.8, the diameter of H is at most $2\beta + O(1) = 2 \lceil \lg 3Cn_2 \rceil + O(1)$. Because n_2 is the number of vertices in the underlying graph, we have $n_2 \leq n(H)$ and therefore the diameter of H is at most $2 \lceil \lg 3Cn(H) \rceil + O(1) = O(\log n(H))$.

It remains to check the size condition on H and k. Because G_1 has the same number of vertices as $G^{\text{NPR}}(\phi)$, Proposition 3.2 implies that the size of G_1 is polynomial in the size of ϕ . Because the underlying graph G_2 is $\mathcal{S}(G_1, \alpha)$ with $\alpha = \max \{ \lg 2t, 4 \lg e(G_1) \}$ and $t \leq 2n(G_1)$, we have that the size of the underlying graph is polynomial in the size of G_1 . Observe that each gadget has size linear in $\beta = \Theta(\log n_2)$. Together with Proposition 5.7, we have that the size of H is polynomial in the size of G_2 . It follows that the size of His polynomial in the size of ϕ . Finally, every gadget's critical distribution size is at most $O(2^{\beta}) = O(n_2)$; together with Proposition 5.7, we have that k is polynomial in n_2 and hence polynomial in n(H).

5.5. **PEBBLING-NUMBER is** Π_2^{P} -complete. After having established Theorem 5.15, it is relatively easy to show that PN is Π_2^{P} -complete.

Theorem 5.16. PEBBLING-NUMBER is Π_2^{P} -complete.

Proof. We have already observed that PN is in $\Pi_2^{\rm P}$. To how that PN is $\Pi_2^{\rm P}$ -hard, we reduce RPN to PN. Let G be a graph with target vertex r, and let $k \ge 0$ be an integer. We produce H and k' so that $\pi(G, r) \le k$ if and only $\pi(H) \le k'$. By Theorem 5.15, our reduction may assume that the diameter d of G is at most $c' \lg n(G)$ for an absolute constant c' and $k = \operatorname{poly}(n(G))$.

We construct H and k' as follows. Let n = n(G) and set $\alpha = \lfloor kn^{c'} \rfloor$. We let H be the graph consisting of α copies of G that share r, so that H - r is α disjoint copies of G - r. We set $k' = \alpha k$. Observe that k' and the size of H are polynomial in the size of G. It remains to show that $\pi(G, r) < k$ if and only if $\pi(H) < k'$.

 (\Longrightarrow) . Suppose $\pi(G, r) \leq k$. Consider a distribution of $k' = \alpha k$ pebbles to H and let u be some target vertex in H. Observe that $d(u, r) \leq d$ and therefore to place a pebble on u, it suffices to show that we can place 2^d pebbles on r. Our strategy is as follows. If there is some copy of G with at least k pebbles, then we arbitrarily select a set S of k pebbles from this copy of G; because $\pi(G, r) \leq k$, we can use these pebbles to place a pebble on r. We repeat this strategy until we are unable to find a copy of G with at least k pebbles. Let s be the number of pebbles we are able to place on r via this strategy. Observe that after executing this strategy s times, at least $k\alpha - ks$ unused pebbles remain in H, and furthermore, if more than $\alpha(k-1)$ unused pebbles remain in H, then some copy of G contains at least k unused pebbles. It follows that

$$k\alpha - ks \leq \alpha(k-1)$$

and therefore $s \ge \alpha/k \ge n^{c'} \ge 2^{c' \lg n} \ge 2^d$.

 (\Leftarrow) . Suppose $\pi(H) \leq k'$, and let p be a distribution of k pebbles to G. Naturally, we define a distribution q of $k' = \alpha k$ pebbles to H by distributing k pebbles in each copy of G according to p. Let D be the signature of a minimum sequence of pebbling moves that places a pebble on r. By the minimum signatures lemma, all edges of D are contained in a single copy of G. It follows that r is reachable in G under p.

6. Conclusions

As we have seen, many graph pebbling problems on unrestricted graphs are computationally difficult. We have seen that REACHABLE and OPTIMAL-PEBBLING-NUMBER are both NP-complete. The authors believe it more likely than not that REACHABLE remains NP-complete even when the graphs are restricted to be planar. However, we have more hope that REACHABLE may fall to P when the graphs are restricted to be

outerplanar. It may be interesting to investigate the computational complexity of these problems when the inputs are restricted to be planar or outerplanar.

We have also seen that PEBBLING-NUMBER is Π_2^{P} -complete, and therefore both NP-hard and coNP-hard. It follows that unless the polynomial hierarchy collapses to the first level, PEBBLING-NUMBER is in neither NP nor coNP. Consequently, given G and k, it is unlikely that we can compute in polynomial time a collection \mathcal{P} of candidate distributions of size k such that if $\pi(G) > k$, then some vertex in G is not reachable from some $p \in \mathcal{P}$ (or else PN would be in NP).

We have shown that COVERABLE and REACHABLE are both NP-complete; however, the computational complexity of these problems diverges when we introduce a universal quantifier over pebble distributions. When we add such a quantifier to COVERABLE, we obtain the problem of determining if $\gamma(G) \leq k$, which is possible in polynomial time [VW, S]. The computational difficulties in COVERABLE are smoothed out by the consideration of all pebble distributions of size k: there is a nice structure to the maximum pebble distributions from which a graph cannot be covered with pebbles. On the other hand, by adding a universal quantifier over all pebble distributions of size k to REACHABLE, we obtain RPN, which asks us to decide if $\pi(G, r) \leq k$. Instead of observing a decrease in the computational complexity, we have stumbled upon a $\Pi_2^{\rm P}$ -complete problem.

We recall that the graph pebbling community has shown a fair deal of interest in developing necessary conditions and sufficient conditions for equality in $\pi(G) = n(G)$. Of course, the ultimate goal is to develop a characterization for when equality holds. We should remark that our hardness result for PEBBLING-NUMBER does not suggest that any such characterization need be complex from a computational point of view. Indeed, our PEBBLING-NUMBER hardness result produces G and k with k > n(G). It may be interesting to explore the complexity of deciding whether $\pi(G) = n(G)$.

Acknowledgements. We thank David Bunde, Jeff Erickson, and Sariel Har-Peled for helpful suggestions throughout the revision process.

References

- [BCCMW] D.P. Bunde, E.W. Chambers, D. Cranston, K. Milans, D.B. West, *Pebbling and Optimal Pebbling in Graphs*, preprint, 2004.
- [C] F.R.K. Chung, Pebbling in hypercubes, SIAM J. Disc. Math, Volume 2 (1989), 467-472.
- [CHH] T.A. Clarke, R.A. Hochberg, and G.H. Hurlbert, *Pebbling in diameter two graphs and products of paths*, J. Graph Th. 25 (1997), 119-128.
- [CCFHPST] B. Crull, T. Cundiff, P. Feltman, G.H. Hurlbert, L. Pudwell, Z. Szaniszlo, and Z. Tuza, *The Cover Pebbling Number of Graphs*, submitted to Discrete Mathematics, 2004.
- [CHKT] A. Czygrinow, G.H. Hurlbert, H. Kierstead, and W.T. Trotter, A note on graph pebbling, Graphs and Combin., 18 (2002) 2, 219-225.
- [H99] G.H. Hurlbert, A survey of graph pebbling, Congressus Numerantium 139 (1999), 41-64.
- [HK] G.H. Hurlbert, H. Kierstead, personal communication.
- [M] D. Moews, *Pebbling graphs*, J. Combin. Theory (B) 55 (1992), 244-252.
- [M-web] D. Moews, *Pebbling graphs*, http://xraysgi.ims.uconn.edu/dmoews/pebbling-graphs.ps
- [PSV] L. Pachter, H.S. Snevily, B. Voxman, On pebbling graphs, Congressus Numerantium 107 (1995) 65-80.
- [P] C.H. Papadimitriou, Computational Complexity, Addison-Wesley, 1994.
- [S] J. Sjöstrand, The Cover Pebbling Theorem, preprint, http://www.arxiv.org/abs/math.CO/0410129, 2004.
- [VW] A. Vuong, M. I. Wyckoff, Conditions for Weighted Cover Pebbling of Graphs, preprint, http://www.arxiv.org/abs/math.CO/0410410, 2004.
- [W] N.G. Watson, The Complexity of Pebbling and Cover Pebbling, preprint, http://arxiv.org/abs/math.CO/0503511, 2005.

Kevin Milans, Computer Science Department, University of Illinois, Urbana, IL, 61801 $E\text{-}mail \ address: milans@uiuc.edu$

Bryan Clark, Physics Department, University of Illinois, Urbana, IL, $6\,180\,1\,$ E-mail address: bkclark@uiuc.edu