

Graph Pebbling: A Blend of Graph Theory, Number Theory and Optimization

Glenn Hurlbert *

Franklin Kenter †

1 Introduction

There are many graph-theoretic subjects that could go under the umbrella of “Moving Things Around on Graphs”. For example, in network optimization, one moves packages from supplies at some vertices to demands at others, with costs accrued per package across edges, attempting to do so most cheaply. In network flow, one tries to maximize the amount of material that can move from a source to a sink, subject to the capacities along edges and the conservation of flow at other vertices along the way. In various versions of pursuit and evasion, cops and robbers take turns moving along edges with the cops trying to capture robbers by landing on them; one tries to minimize the number of cops necessary to always capture the robbers. Zero forcing, black pebbling, chip firing, black and white pebbling, graph pegging, and rigid pebbling are other examples, all with different rules for moving, different objectives, and different constraints and/or costs. While some of these topics model real world problems, others model pure mathematical problems such as space-time tradeoffs in computational complexity, efficient matrix storage during Cholesky factorization, rigid structures in computational geometry, and matrix rank and nullity computation. The graph pebbling model we discuss here is used to translate a number-theoretic problem into graph theory. It has also developed into a rich subject in its own right.

In this article, we present a range of paradigms within this field and, for each one, share some of the major results, conjectures, and open problems. Many of these problems are accessible to those both at many levels, from undergraduate (e.g. [7, 8, 9, 19]) to seasoned emeritus, and from various backgrounds, from algebra and graph theory to probability and discrete optimization. In fact, many well-known professionals (e.g. Bukh, Elledge, Gibbons, Herscovici, Moews, Pachter, Pudwell, Tomova, Wierman, Xue, Yerger, others) have produced great graph pebbling results during their time as undergraduates. For interested readers, the two resources [12, 13] contain references for most of the results not cited below.

2 Number-Theoretic Origins

2.1 Zero Sum Subsequences

The story of graph pebbling begins with Erdős and zero-sum problems. Zero-sum problems are combinatorial problems in the context of finite abelian groups that determine how large a set (or sequence) of elements must be such that some subset (or subsequence) with prescribed constraints has a zero sum.

Theorem 1 (Erdős-Ginzburg-Ziv, 1961). *Let n be a positive integer. Any sequence of length $2n - 1$ of elements from \mathbb{Z}_n has a subsequence of length exactly n that sums to $0 \pmod n$.*

This is best possible; consider the sequence $n - 1$ ‘0’s and $n - 1$ ‘1’s, which has no nonempty subsequence of length n that sums to $0 \pmod n$.

More generally, one can ask all sorts of questions regarding the length of a sequence before the sum of some (nonempty) subsequence has a specified property. For instance, if S is a sequence of $|\Gamma|$ elements of a group Γ , then the pigeonhole principle implies that there is nonempty zero sum subsequence. The minimum

*Department of Mathematics and Applied Mathematics, Virginia Commonwealth University, USA

†Mathematics Department, United States Naval Academy, USA

such length for which a zero sum subsequence is guaranteed is called the *Davenport constant* of Γ . For \mathbb{Z}_n , n is best possible (consider a sequence of $n - 1$ ‘1’s), so the Davenport constant of \mathbb{Z}_n is n .

Erdős and Lemke conjectured that there is a nonempty zero sum subsequence T of $S = (s_1, \dots, s_n)$ with the additional condition that $\sum_{t \in T} t \leq \text{lcm}(s_1, \dots, s_n, n)$. A stronger condition is that $\sum_{t \in T} \text{gcd}(t, n) \leq n$, which can be written as $\sum_{t \in T} \frac{1}{|t|} \leq 1$, where we write $|g|$ for the order of an element g in a group Γ (in this case \mathbb{Z}_n). This sum is often referred to as the *cross number* of the subsequence T , which plays a vital role in factorization theory. Kleitman and Lemke answered this stronger question in the affirmative.

Theorem 2 (Lemke-Kleitman, 1989). *Let n be a positive integer. Let S be a sequence of \mathbb{Z}_n with length n . Then there is some subsequence $T = (t_1, t_2, \dots, t_q)$ of S such that both*

$$\sum_{t \in T} t \equiv 0 \pmod{n} \quad \text{and} \quad \sum_{t \in T} \frac{1}{|t|} \leq 1.$$

The natural extension of their work is to arbitrary groups, which remains open for nonabelian Γ .

Conjecture 3 (Lemke-Kleitman, 1989). *Let Γ be a finite additive group. Let S be a sequence of Γ with length $|\Gamma|$. Then there is some subsequence $T = (t_1, t_2, \dots, t_q)$ of S such that both*

$$\sum_{t \in T} t = 0_{\Gamma} \quad \text{and} \quad \sum_{t \in T} \frac{1}{|t|} \leq 1.$$

2.2 Chung’s Proof of the Lemke-Kleitman Theorem

The original proof of the Lemke-Kleitman Theorem relied on an inductive argument based on a multitude of possible interactions between elements. Lagarias and Saks proposed a simplified approach to the Erdős-Lemke problem, and Chung was the first one to introduce it to the literature [6]. The main concept is to combine elements with like prime factorizations. Let $n = p_1^{a_1} p_2^{a_2} \dots p_q^{a_q}$ be the prime factorization of n . For each element s_i in the sequence S , write $s_i = p_1^{a_{i,1}} p_2^{a_{i,2}} \dots p_q^{a_{i,q}} R_i$ where R_i is relatively prime to n .

Choose any exponents (b_1, \dots, b_q) and let T be a subsequence of S containing elements of the form $p_1^{b_1} p_2^{b_2} \dots p_q^{b_q} R$ (where R is not necessarily relatively prime to n). If, for some j , $|T| \geq p_j$ then there is subsequence T' of T , of size at most p_j , such that the sum of the elements in T' necessarily has the form $p_1^{b_1} p_2^{b_2} \dots p_j^{b_j+1} \dots p_q^{b_q} R$. Furthermore, if each $t \in T$ satisfies

$$1/|t| \leq p_1^{b_1-a_1} p_2^{b_2-a_2} \dots p_j^{b_j-a_j} \dots p_q^{b_q-a_q}$$

then

$$\sum_{t \in T'} 1/|t| \leq p_1^{b_1-a_1} p_2^{b_2-a_2} \dots p_j^{b_j+1-a_j} \dots p_q^{b_q-a_q}.$$

Thereafter, we will treat T' as a single entity, where the value of T' is $\sum_{t \in T'} t$, and the notation $1/|T'|$ is taken to denote $\sum_{t \in T'} 1/|t|$. Note that these definitions are applied recursively, as an element $t \in T'$ may itself be a set, rather than a number. By iteratively and strategically applying such groupings, one can find a zero-sum subsequence of S that also obeys the cross number condition. This is exactly how the Chung’s proof of the Lemke-Kleitman Theorem works. In this sense, her proof can be viewed as an abstract puzzle game. That game is graph pebbling!

The translation of the Erdős-Lemke problem into graph pebbling is not trivial. The essence of the proof is to (1) build the divisor lattice $L(n)$, (2) view the given numbers as pebbles, placed naturally in $L(n)$ according to common factors with n , (3) observe that each pebble obeys local versions of the two required conditions, (4) model certain movements of the pebbles in the lattice by “pebbling steps” (to be defined later), (5) show that those steps preserve the two local conditions as illustrated above, (6) observe that any pebble that reaches the bottom represents a solution to Erdős-Lemke problem, and finally (7) prove that it is always possible to move a pebble to the minimal vertex via pebbling steps, given *any* configuration of n pebbles. A detailed translation is given by Elledge and Hurlbert, who gave a Chung-like pebbling proof of the Lemke-Kleitman conjecture for abelian groups (originally resolved by Geroldinger).




Figure 1: The upside-down divisor lattice of $n = 1,321,375 = 5^3 \cdot 11^1 \cdot 31^2$, with the number $p = 2,984,850 = 5^2 \cdot 11^1 \cdot 31^0 \cdot 2 \cdot 3^4 \cdot 67$ placed as a pebble at $(2, 1, 0)$.

We now illustrate the pebbling movement of this approach with a specific example as seen in Figure 1. Consider a sequence of 5 pebbles that are placed at $(2, 1, 0)$, such as $T = (550, 1925, 1100, 3300, 1100)$. Because these are all multiples of $5^2 \cdot 11^1 \cdot 31^0 = 275$, it is simpler to represent them just as those multiples, and simpler still to reduce them modulo 5, becoming the sequence $(2, 2, 4, 2, 4)$. Now it is easier to spot a zero-sum in \mathbb{Z}_5 , such the final three numbers $(4, 2, 4)$. Then we can move, as a pebbling step, the corresponding subsequence $T' = (1100, 3300, 1100)$ down the red edge to $(3, 1, 0)$ as a single “superpebble”, labeled by the sequence T' . Notice that $\sum_{t \in T'} t = 5500 = 5^3 \cdot 11^1 \cdot 31^0 \cdot 2^2$, and that $\sum_{t \in T'} 1/|t| = 3(5^{-1} \cdot 11^0 \cdot 31^{-2}) \leq 5(5^{-1} \cdot 11^0 \cdot 31^{-2}) = 5^0 \cdot 11^0 \cdot 31^{-2}$. This is the essence of how pebbling models both sums. Any superpebble that reaches the bottom will be a zero-sum with cross number at most 1.

The final stage (7) above is handled by reducing to the square-free case. That is, if $n = p_1^{a_1} p_2^{a_2} \dots p_q^{a_q}$, we reduce the proof on $L(n)$ to that on $L(m)$, where m is any product of d distinct primes with $d = \sum a_i$, often referred to as Q^d , the lattice subsets of $\{1, \dots, d\}$. One can see that $L(n)$ is the product of q paths, having lengths a_1, \dots, a_q , while Q^d is the product of d paths, each of length 1. Because $L(n)$ is a sublattice of Q^d (see Figure 2) — in fact, it is a *retract*: a homomorphic image of Q^d — homomorphic images of pebbling steps in Q^d are pebbling steps in $L(n)$. Chung then uses the stronger recursive nature of Q^d to solve the problem there more easily. We will illustrate her technique in Section 3.4.

3 Lagarias-Saks Pebbling

We will assume basic familiarity with graphs as sets of vertices and edges (unordered pairs of vertices), and use the notation $n(G)$ for the number of vertices of a graph G . All graphs we consider are connected. Common notation for graphs on n vertices includes the complete graph K_n in which every pair of vertices is adjacent, the path P_n in which only consecutive vertices are adjacent, and the cycle C_n which adds an edge between the endpoints of P_n . For two vertices u and v , the *distance* function $\text{dist}(u, v)$ measures the fewest number of edges in a path from u to v . The *eccentricity* of a vertex v of G is defined $\text{ecc}_G(v) = \max_u \text{dist}(u, v)$, while the *diameter* of a graph G is defined $\text{diam}(G) = \max_v \text{ecc}_G(v)$. Sometimes we drop the subscript of G when the context is clear. We write $\lg(\cdot)$ for the base 2 logarithm. Additionally, we write $f \in O(g)$ (resp.




Figure 2: Embedding the divisor lattice $L(p_1^2 \cdot p_2^2)$ into Q^4 as a retract. The thick edges display $L(n)$. Then the thin edges can be mapped to the thick ones (of the same color) in an adjacency-preserving manner.

$\Omega(g)$ if $|f(x)| < Cg(x)$ (resp. $> C$) for some C and all large enough x ; and $\Theta(g) = \Omega(g) \cap O(g)$. We also use $f \ll g$ to mean that $f(x)/g(x) \rightarrow 0$ as $x \rightarrow \infty$.

3.1 Pebbling Basics

A *configuration* C of pebbles on the vertices of a connected graph G is a function $C : V(G) \rightarrow \mathbb{N}$ (the nonnegative integers), so that $C(v)$ counts the number of pebbles placed on the vertex v . We write $|C|$ for the *size* $\sum_v C(v)$ of C ; i.e. the number of pebbles in the configuration. A *pebbling step* from a vertex u to one of its neighbors v (denoted $u \rightarrow v$) reduces $C(u)$ by two and increases $C(v)$ by one. Given a specified *target* vertex r we say that C is *t -fold r -solvable* if some sequence of pebbling steps places t pebbles on r . Conversely, if no such steps exist, then C is *r -unsolvable*. We are concerned with determining $\pi_t(G, r)$, the minimum positive integer m such that every configuration of size m on the vertices of G is t -fold r -solvable. The invariant is well-defined since, for large enough t , the pigeonhole principle guarantees that some vertex will contain at least $2^{\text{ecc}_G(r)}$ pebbles, which can by itself solve r . The *t -fold pebbling number* of G is defined to be $\pi_t(G) = \max_{r \in V(G)} \pi_t(G, r)$. When $t = 1$, we simply write $\pi(G)$, which is the *pebbling number* of G .

The configuration with a single pebble on every vertex except the target shows that $\pi(G) \geq n(G)$, the number of vertices of the graph G , while the configuration with $2^{\text{ecc}_G(r)} - 1$ pebbles on the farthest vertex from r , and no pebbles elsewhere, shows that $\pi(G) \geq 2^{\text{diam}(G)}$ when r is chosen to have $\text{ecc}_G(r) = \text{diam}(G)$.

One can also view Q^d as the graph on all binary d -tuples, pairs of which that differ in exactly one coordinate are joined by an edge. In this guise, Q^d is often referred to as the *d -dimensional hypercube*. As part of her work on the Lemke-Kleitman Theorem, Chung proves the following.

Theorem 4 (Chung [6]). $\pi(Q^d) = 2^d$.

In particular, both of the previous bounds are simultaneously tight, as $n(Q^d) = 2^d = 2^{\text{diam}(Q^d)}$.

To provide an alternative solution to the Erdős-Lemke problem, Chung actually showed more. Given a positive integer n , let its factorization be $p_1^{a_1} \cdots p_q^{a_q}$ and define $d = \sum_{i=1}^q a_i$. Alter the pebbling step rule so that, in a_i of the dimensions of Q^d , p_i pebbles (instead of 2) must be removed from a vertex to move one pebble in those dimensions. Chung showed in this case that the pebbling number of Q^d with these generalized pebbling steps is n , and so the Lemke-Kleitman Theorem holds.

3.2 Class 0 graphs: When does $\pi(G) = n(G)$?

Returning to the two-to-one pebbling rule, graphs G that, like Q^d , have $\pi(G) = n(G)$ have come to be known as *Class 0*. The terminology comes from a lovely theorem of Pachter, Snevily and Voxman, which states that if $\text{diam}(G) = 2$ then $\pi(G) \leq n + 1$; thus there are two classes of diameter two graphs — Class 0 and Class 1. If $G - v$ is disconnected for some vertex v then $\pi(G) > n$. Let G_1 and G_2 be different connected components of $G - v$ with $u_i \in G_i$, and define the configuration C by $C(u_1) = C(v) = 0$, $C(u_2) = 3$, and




Figure 3: The pyramid graph, with r -unsolvable configuration shown.

$C(w) = 1$ for all other vertices w . Then, C is u_1 -unsolvable of size n . The minimum number of vertices whose removal disconnects a graph is called its *connectivity*. Thus, Class 0 graphs are 2-connected.

There are other properties that Class 0 graphs must have. For example, denote by $N_a[v]$ the set of vertices of distance at most a from vertex v . Cranston et al. [7] observe that if G is Class 0 and u and v are vertices with $N_a[u] \cap N_b[v] = \emptyset$, for some a and b , then $|N_a[u] \cup N_b[v]| \geq 2^{a+b+1}$ — otherwise the configuration with $2^{a+b+1} - 1$ pebbles on v , no pebbles on other vertices of $N_a[u] \cup N_b[v]$, and 1 pebble on each remaining vertex, is u -unsolvable of size at least n . They use this *Neighborhood Lemma* to prove that every Class 0 graph with $n \geq 4$ vertices has at least $(5n - 11)/3$ edges. At present, the Class 0 graph with the fewest edges known is the *wheel*, a cycle with a central vertex adjacent to every cycle vertex, having $2n - 2$ edges. Are there Class 0 graphs with fewer edges?

The family of 2-connected Class 1 diameter-2 graphs were later characterized by Clarke, Hochberg and Hurlbert, the main property of which is the induced containment of the (Class 1) *pyramid* shown in Figure 3. The other important property is that any every vertex not in the pyramid lives geometrically inside one of its triangles, and vertices from different triangle interiors are not adjacent. In a probabilistic sense, almost all graphs have diameter two — think of the probability that some two vertices don't have a common neighbor. Now consider how unlikely it is to fit into this restrictive pyramid structure: at best $O((1/8)^n)$ of all graphs are like this since, if there is a pyramid, each of the $n - 6$ non-pyramid vertices have at least 3 non-neighbors. Thus, almost all graphs are Class 0.

One can say more. As shown by Postle et al. [20], graphs of diameter three can have pebbling numbers as high as $3n/2 + 2$, but no more, while graphs of diameter four have pebbling numbers at most $(3/2)n + \Theta(1)$. The techniques used to achieve these results include “discharging” which is the main technique used in the proof of the celebrated Four-Color Theorem. A higher diameter can push pebbling numbers to the exponential range (e.g. $\pi(P_n) = 2^{n-1}$). In general, Postle showed that a graph with diameter d has pebbling number at most

$$\frac{2^{\lceil \frac{d}{2} \rceil} - 1}{\lceil \frac{d}{2} \rceil} n + O(\sqrt{n \ln n}).$$

On the other hand, high connectivity (ensuring many paths to the target) can keep pebbling numbers low. Czygrinow et al. [10] proved that for all d , there is a $k(d)$ such that every graph with diameter d and connectivity at least $k(d)$ is Class 0. The Erdős-Renyi random graph model places edges independently with probability p between pairs of vertices (the counting above is for the model $p = 1/2$). The authors proved that $k(d) \leq 2^{2d+3}$ (it was shown by Clarke et al. to be at least $2^d/d$), and used the result to prove that a random graph with edge probability $p \gg (n \lg n)^{1/d}/n$ is almost surely Class 0. The proof used known theorems about the thresholds for having connectivity at least k and diameter at most d . It is interesting how close this value is to the connectivity threshold of $\lg n/n$, yet there is plenty of room for improvement in the upper and lower bounds for $k(d)$.

One nice application of the diameter-connectivity theorem is to Kneser graphs, known for their importance in many subareas of graph theory and combinatorics, including graph coloring, graph embedding, and extremal set theory, among others. For $m \geq 2h + 1$, the Kneser graph $K(m, h)$ has all h -subsets of $[m] = \{1, \dots, m\}$ for vertices, with edges joining disjoint pairs. $K(m, 1)$ is the complete graph K_m , and the famous Petersen graph $P = K(5, 2)$ in Figure 4 is the smallest example for $h \geq 2$. Readers should find it an enjoyable puzzle to prove that the Petersen graph is Class 0. (There are at least five distinct proofs in




Figure 4: The Petersen graph $P = K(5, 2)$.

the literature, one of which follows from the diameter two classification above; two others appear below.) In fact, it was shown in [12] that $K(m, 2)$ is Class 0 for all $m \geq 5$. Furthermore, with the diameter-connectivity result in hand, they prove that, for any constant $c > 0$, there is an integer h_0 such that, for all $h > h_0$ and $s \geq c(h/\lg h)^{1/2}$, $K(2h + s, h)$ is Class 0. With this evidence, one cannot help but believe that $K(m, t)$ is Class 0 for all $m \geq 2h + 1$, and that some clever idea to prove it lies in wait. The subject could use more Class 0 sufficiency conditions.



Figure 5: A configuration of weight $127/128$ on the path P_8 with target r .

3.3 Weight Functions: An Optimization Approach

Along these lines, there is an opportunity to use weight functions. One can think of pebbles sitting on a target as having weight 1, on a target's neighbors as having weight $1/2$, ..., and sitting on a vertex at distance d from the target as having weight $1/2^d$. The *weight* of a configuration is then the sum of the weights of its pebbles (see Figure 5). It is not difficult to prove, for example, that a configuration on a path can solve the target at one of its endpoints if and only if its weight is at least one. By rescaling to make all the weights integral, we can rephrase the characterization as: a configuration on a path cannot solve the target at one of its endpoints if and only if its weight is less than 2^{n-1} ; i.e. at most $2^{n-1} - 1$. Interestingly, this bound equals the sum of the weights.

This can be generalized at the cost of equivalence, and the notion leads us to techniques involving linear and integer optimization. Suppose that T is a tree, with one of its leaves r as target, and that $\text{ecc}_T(r) = d$. Define the *weight function* w_T by $w_T(v) = 2^{d-i}$, where $i = \text{dist}_T(v, r)$. The pair (T, w_T) is called a *basic strategy*. The Weight Function Lemma [14] says that if C is an r -unsolvable configuration on T then

$$\sum_{v \in T} w_T(v)C(v) \leq \sum_{v \in T} w_T(v). \tag{1}$$

One can extend this to all graphs as follows. Given a target r in a graph G , consider any tree T in G with r as a leaf. Extend w_T to all of G by setting $w_T(v) = 0$ when $v \notin T$; then the Weight Function Lemma still holds. The collection of all inequalities from basic strategies gives rise to an integer optimization problem by maximizing $z_r = \sum_{v \in G} C(v)$ (the size of C) over these constraints. If z_r^* is the optimum value then it shows that every r -unsolvable configuration has size at most z_r^* ; in other words, $\pi(G, r) \leq z_r^* + 1$. This gives rise to the following integer program formulation of this graph pebbling upper bound:

$$\begin{aligned}
\max \quad & 1 + \sum_v C(v) \\
\text{s.t.} \quad & \sum_{v \in T} w_T(v) C(v) \leq \sum_{v \in T} w_T(v) \quad \forall \text{ basic strategies } (T, w_T).
\end{aligned}$$

There are some drawbacks to the approach, as you might imagine. Of course, integer optimization is typically exponential in the input size, but the input size (the number of such trees) is typically exponential (or worse!) in the number of vertices. Furthermore, consider the *unsolvability polytope* of an n -vertex graph G with target r : the convex hull of all r -unsolvable configurations on G (viewing configurations as non-negative vectors in \mathbb{R}^n). The Weight Function Lemma implies that the feasible polytope of the linear optimization problem contains the unsolvability polytope — but it can be a very slack containment. To tighten that slack, Cranston et al. [7] developed additional constraints based on certain non-tree subgraphs called lollipops (a path with a cycle attached to an endpoint). For instance, the tree strategies only yield $\pi(Q^3) \leq 9$, while the addition of lollipops proves $\pi(Q^3) = 8$.




Figure 6: A [tree constraint](#) in the Petersen graph P .

In practice, however, one can have surprisingly good success by restricting the input to a random collection of polynomially many breadth-first subtrees of depth at most $\text{ecc}_G(r) + 1$, for example, and then stop with just linear optimization and the inequality $\pi(G, r) \leq \lfloor z_r^* \rfloor + 1$. The resulting approximation can be very close, even exact, especially when the diameter of the graph is not large. We can illustrate this on the Petersen graph P in Figure 6. The tree T_{35} shown in blue corresponds (using the labeling from Figure 4) to the inequality

$$4C(v_{35}) + 2C(v_{14}) + 2C(v_{24}) + C(v_{23}) + C(v_{25}) + C(v_{13}) + C(v_{15}) \leq 12.$$

Using the symmetries yields two other trees T_{34} and T_{45} . These produce similar constraints, and the sum of all three constraints reveals that

$$4C(v_{34}) + 4C(v_{35}) + 4C(v_{45}) + 4C(v_{14}) + 4C(v_{24}) + 4C(v_{23}) + 4C(v_{25}) + 4C(v_{13}) + 4C(v_{15}) \leq 36, \quad (2)$$

so $|C| \leq 9$, and hence $\pi(P, v_{12}) \leq 10$. Combined with the lower bound from having 10 vertices, this shows (with the symmetry that all vertices look the same) that the Petersen graph is Class 0. This example also illustrates the following lemma.

Lemma 5 (Uniform Covering Lemma [14]). *Let r be a target vertex in a graph G and \mathcal{T} be a set of r -strategies. If there is some m such that $\sum_{T \in \mathcal{T}} T(v) = m$ for all $v \neq r$ in G then G is Class 0 at r .*

This holds because the right hand side of inequality (1) is the sum of the coefficients, so the same is true of any sum of such weight functions. Therefore, in a uniform sum such as inequality (2), the right hand side, after dividing by m , equals $n - 1$, the number of non-target vertices. The left hand side, with uniform coefficients, always becomes the size of C after this division.

In some cases the method of weight functions offers the simplest proof of a graph's pebbling number (or the best upper bound of it). The set of strategies used in such a proof is called a *certificate*. Changing

the weighting system by reducing weights by more than half as we traverse the edges of a tree, proceeding away from the target, yields *non-basic* strategies — it turns out that these are conic (i.e. non-negative linear) combinations of basic strategies, and so satisfy the Weight Function Lemma as well. Thus, non-basic strategies do not need to be included in the optimization constraints, but they can be used in a certificate. From the geometrical point of view of the unsolvability polytope, certificates will involve at most $n - 1$ strategies. However, in practice, most certificates have size $\deg_G(r)$ because two strategies using the same neighbor of r can often be combined. Hence, finding certificates for a particular graph can even be quicker to do by hand than by computer, especially when symmetry is present, and finding them for an infinite family of graphs is really a hands-only activity.

A good example of this is powers of cycles [14]. For a graph G , define the k^{th} power of G , denoted $G^{(k)}$, by adding the edge uv for every pair of vertices satisfying $\text{dist}_G(u, v) \leq k$. The *pebbling exponent* of G , denoted $e_\pi(G)$, is defined to be the minimum k for which $G^{(k)}$ is Class 0. For instance, $G^{(\text{diam}(G))} = K_n$ for every G on n vertices, and so $e_\pi(G) \leq \text{diam}(G)$. Because Class 0 graphs must have $2^{\text{diam}(G)} \leq n(G)$, we need that $\text{diam}(G^{(e_\pi(G))}) \leq \lg n(G)$. This provides the lower bound on the result for n -vertex cycles that

$$\frac{n/2}{\lg n} \leq e_\pi(C_n) \leq \frac{n/2}{\lg n - \lg \lg n},$$

where the upper bound comes from a certificate of carefully constructed strategies.

The following problem is an interesting question from [14] regarding the approximability of graph pebbling with weight functions. Let $\mathcal{T} = \mathcal{T}(G, r)$ be the set of all tree r -strategies in a graph G . Is the resulting linear optimization solution $z_r^*(\mathcal{T})$ a reasonable approximation to $\pi(G, r)$? That is, does there exist some constant c such that $z_r^*(\mathcal{T}) \leq c\pi(G, r)$? Is $c = 2$?

Optimization perspectives of graph pebbling are an active area of research. For instance, Kenter et al. [15] presents optimization approaches specific to graph products. In general, describing the actual facets of the unsolvability polytopes of these approaches also remains illusive.

3.4 Graham's Conjecture: Pebbling on Graph Products

Chung's proof that $\pi(Q^d) = 2^d$ is inductive, but for the induction to work, it requires also proving that Q^d has the *2-Pebbling Property (2PP)*: two pebbles can be placed on any target of G from any configuration C of size at least $2\pi(G) - s(C) + 1$, where $\text{supp}(C)$ is the set of vertices having at least one pebble (the *support* of C) and $s(C) = |\text{supp}(C)|$.

As a toy example of Chung's proof technique, we prove again that the Petersen graph P is Class 0, mimicking her three cases. Observe that P splits into two 5-cycles A and B with a matching between them, suppose by symmetry that $r \in A$, and let C be a configuration of size 10. If A contains at least 5 pebbles then we are done because $\pi(A) = 5$, so assume not. Let r' be the neighbor of r in B . If r' contains a pebble then the other five pebbles in B can place a second pebble on r' , subsequently solving r , so presume otherwise. If A contains exactly 4 (resp. 3) pebbles then the 6 (resp. 7) pebbles in B can move at least 1 (resp. 2) pebbles across to A , thus solving r from the resulting 5. If A contains at most 2 pebbles then B has at least 8, from which we can place two pebbles on r' , then one on r .

Because Chung needed to prove inductively that the cube had both pebbling number n and the 2PP, she needed also to prove the 2PP in three similarly partitioned cases.

As usual, one might look for a simpler or more direct proof of Chung's result. In this case we observe that Q^d is a d -dimensional product of edges. The *Cartesian product* of two graphs G and H is denoted $G \square H$. Its vertices are the pairs (u, v) with $u \in V(G)$ and $v \in V(H)$, and two vertices (u_1, v_1) and (u_2, v_2) form an edge when $u_1 = u_2$ and $v_1 v_2$ is an edge of H and when $v_1 = v_2$ and $u_1 u_2$ is an edge of G . (Observe the notation of the empty box, which represents the equality $K_2 \square K_2 = C_4$.) For example, Figure 7 shows the case when G is a 4-vertex star and H is a triangle. Now one can see that $Q^d = Q^2 \square Q^{d-1}$. Thus, a simple proof would follow inductively from the relation that $\pi(Q^d) \leq \pi(Q^2)\pi(Q^{d-1})$, since $\pi(Q^2) = 2$. This led Graham to make the following conjecture, which many consider the holy grail of the subject.

Conjecture 6 (Graham's Conjecture [6]). *For all graphs G and H we have $\pi(G \square H) \leq \pi(G)\pi(H)$.*

There has been a significant amount of work verifying Graham's conjecture for a wide array of cases, including when G is a product of complete graphs and H has the 2PP [6], when G and H are products of




Figure 7: Graphs $G = S_4$ (left), $H = K_3$ (center), and $G \square H = S_4 \square K_3$ (right).




Figure 8: The Lemke graph L , with a 2-fold r -unsolvable configuration C .

trees [19], when G and H are products of cycles (Herscovici), when G and H are complete bipartite graphs on at least 15 vertices per part (Feng and Kim), when G is a product of a tree and H has the 2PP (Herscovici) and many more.

The proofs of these results (and by “products of” we mean to include a graph by itself, as a trivial product) all use that one of the graphs has the 2PP. This led to investigations about whether or not all graphs have the 2PP — including products of complete graphs [6] and trees [6], among others. Lemke ([6]) found a graph with the smallest number of vertices not having the 2PP. It is not difficult to show that the *Lemke graph* L in Figure 8 is Class 0 (try it with certificates if you like). However, the configuration C shown has size $2\pi(L) - s(C) + 1 = 12$, and yet two pebbles cannot be moved to the target r . Indeed, define the weight $w(v) = 2^{-\text{dist}(v,r)}$ of a vertex v , and corresponding weight $w(C) = \sum_{v \in G} w(v)C(v)$ of a configuration C . Then, since the weight of the configuration C shown is 2, a 2-fold r -solution σ cannot afford to lose any weight. Thus, every pebble must be used in σ , and σ must be *greedy* — every pebbling step moves closer to the target. But this forces three pebbles to b and only one to a , thereby placing only one pebble on r . Any graph without the 2PP has come to be known as a *Lemke graph* (as opposed to *the* Lemke graph), and these are studied in their own right (see, for instance, [9]). Indeed, there are 22 Lemke graphs with the minimum of 8 vertices; each is a subgraph of one having 17 edges, and a unique one (not L !) has the minimum of 12 edges. Additionally, several infinite families of Lemke graphs are known.

One result that avoids the necessity of the 2PP is that of Czygrinow and Hurlbert which states that if G and H are graphs on at most n vertices each, the minimum degree of a vertex in G or H is δ , and that $\delta \geq 2^{(12n/\delta)+15}$, then $G \square H$ is Class 0. From this it follows that $\pi(G \square H) = n(G \square H) = n(G)n(H) \leq \pi(G)\pi(H)$. Essentially, this says that Graham’s conjecture holds for graphs of minimum degree at least $cn/\lg n$ for some c . It’s the density of such graphs that makes their product satisfy the diameter-connectivity theorem. Repeated applications of this theorem die out eventually, however, since the degrees grow additively while the number of vertices grows multiplicatively.

We see in this argument that the truth of Graham’s conjecture would imply that Class 0 graphs would be closed under Cartesian products. Hence verifying the conjecture for Class 0 graphs would be a very interesting result. Unfortunately, little has been achieved in this direction, possibly because of the lack of

sufficient conditions for being Class 0.

Another way to avoid using the 2PP is given by Herscovici . Define a graph G to have the *path property* if $\pi(G \square P_n) \leq \pi(G)\pi(P_n)$ for all n . Any graph without the path property is a counterexample to Graham’s conjecture, so this somewhat simple instance seems a good place to start. Thus, once proven, one could then prove Graham’s conjecture by proving that it holds for graphs with the path property. This is a promising line of reasoning.

On the other hand, if you like to play devil’s advocate, since most proofs do use the 2PP, it seems reasonable to wonder if Graham’s conjecture might fail for $L^2 = L \square L$. Because L is Class 0, we should have $\pi(L^2) = 64$; but is it? In the tradition of Paul Erdős, the author of [14] offers \$64 for the resolution of this question ¹.

Evidence in favor of Graham’s conjecture, in particular that the 2PP should not be relevant, is given by Gao and Yin. Snevily and Foster defined an infinite family of graphs $\{L_i\}_{i \geq 1}$, which Gao and Yin proved are all Lemke. Wang produced another sequence $\{L'_i\}_{i \geq 1}$ of Lemke graphs. They showed that, for every $L_i \subseteq G \subseteq L'_i$, if H is a complete graph or tree then $\pi(G \square H) \leq \pi(G)\pi(H)$.

Additionally, Pleanmani has shown in that, for any G that is not Class 0, if $n \geq 2(\pi(G)-1)/(\pi(G)-n(G))$ then $\pi(G \square K_n) \leq \pi(G)\pi(K_n)$. They proved an analogous result in for large enough complete bipartite graphs as well.

Work has begun on approximations to Graham’s conjecture, as Asplund et al. [2] proved that $\pi(G \square H) \leq 2\pi(G)\pi(H)$ always holds. They actually proved the stronger result that $\pi(G \square H) \leq (\pi(G) + n(G))\pi(H)$. This is no better a bound for Class 0 graphs, but means that

$$\pi(G \square H) \leq (1 + \epsilon)\pi(G)\pi(H) \text{ whenever } n(G) \leq \epsilon\pi(G),$$

such as for graphs having $\text{diam}(G) \geq \lg(n/\epsilon)$. Can better approximations than 2 be found for all graphs?

Another variation on the problem is to change the graph product. The authors of [2] consider several, among them the *strong* product $G \boxtimes H$, which adds to $G \square H$ edges between (u_1, v_1) and (u_2, v_2) when u_1u_2 is an edge of G and v_1v_2 is an edge of H . (Observe the notation of the full box, which represents the equality $K_2 \boxtimes K_2 = K_4$.) Since $G \square H \subseteq G \boxtimes H$, we have $\pi(G \boxtimes H) \leq \pi(G \square H)$. Regarding a Graham-type inequality, they prove that every G and H satisfy $\pi(G \boxtimes H) \leq \frac{3}{2}(\pi(G) + 1)(\pi(H) + 1)$.

3.5 Computational Complexity: How Hard is Graph Pebbling?

The problem of computing $\pi(L^2)$, mentioned above, begs the following questions.

1. How hard is it to decide if a particular configuration C on G solves a particular target r ?
2. If we know that C solves r , how hard is it to display an actual solution σ ?
3. And, of course, how hard is it to compute $\pi(G)$?

For a quick refresher in complexity theory, P is the set of problems that can be solved in polynomial time (polynomial as measured by the size of the input, in our case, $n(G)$), NP is the class of problems that can be verified in polynomial time, and EXPTIME is the set of problems that can be solved in exponential time. The set NP-complete (resp. EXPTIME-complete) is the class of problems for which all other NP (resp. EXPTIME) problems can be reduced to.

For Question 1, it was proved by both Hurlbert and Kierstead and by Milans and Clark [18] that the problem of deciding if C solves r , for a given instance of configuration C on a graph G with target vertex r , is NP-complete. The former pair reduced the well known NP-complete problem of deciding if a 4-uniform hypergraph (a collection of subsets of vertices, each of size 4; e.g. a graph is a 2-uniform hypergraph) has a perfect matching (a set of edges that partitions its vertices) to this problem of pebbling solvability. The short proof illustrates another nice weight argument, so it’s worth presenting here.

Let \mathcal{H} be 4-uniform hypergraph on 2^{m+2} vertices, with edges $E(\mathcal{H}) = \{e_1, \dots, e_k\}$ (it is not difficult to show that prescribing this number of vertices is not a loss of generality). Define the *pebbling graph* $G = G_{\mathcal{H}}$ as follows. Its vertices are given by $V(G) = V(H) \cup E(H) \cup \{u_1, \dots, u_k\} \cup \{r, w_1, \dots, w_m\}$. The edges $E(G)$ include ve_i for every $v \in e_i$, as well as the paths $w_mu_ie_i$ for all $i \leq k$, and the path $rw_1 \dots w_m$ (see Figure

¹This is not to be confused with the \$64,000 question!




Figure 9: The pebbling graph $G_{\mathcal{H}}$ with configuration $C_{\mathcal{H}}$.

9). Now define the configuration $C = C_{\mathcal{H}}$ by $C(v) = 2$ for all $v \in V(\mathcal{H})$ and $C(v) = 0$ otherwise. We then argue that C solves r if and only if \mathcal{H} has a perfect matching.

Define the weight of a configuration as above. Notice that the weight of a configuration never increases via pebbling steps, and is preserved if and only if the pebbling step is greedy. Therefore, since a configuration with a pebble on the target r has weight at least 1, any r -solvable configuration must have weight at least 1. In particular, since this configuration on $G_{\mathcal{H}}$ has weight 1, it can only solve r through greedy steps, with all pebbles used up in the process. In other words, all pebbling steps in Figure 9 would need to move from right to left.

This means two things. First, each e_i can receive at most 4 pebbles, because e_i has exactly four edges to $\{v_1, \dots, v_n\}$. Second, each pebble that w_m receives from some e_i comes from exactly four pebbles at e_i . Therefore, the number of pebbles that reach e_i from the right is either 0 or 4. In addition, no two e_i and $e_{i'}$ that receive pebbles share a common neighbor v_j , since v_j can send a pebble to only one of them; that is, as edges of \mathcal{H} , $e_i \cap e_{i'} = \emptyset$. Now, it takes 2^m pebbles on w_m to reach r , and so it must receive them from a set of 2^m distinct e_i s — it is that set of e_i s that forms a perfect matching in \mathcal{H} .

It is easier to see now that a perfect matching in \mathcal{H} does yield an r -solution, which completes the proof.

As is customary with NP-complete problems, one typically searches for classes of graphs in which the problem is polynomial, and here we have some nice results. For example, the proof of NP-completeness by Milans and Clark uses reduction from 3SAT, a prototypical NP-complete problem, and shows that the problem remains NP-complete for bipartite graphs, even those having maximum degree 3. Similarly, Cusack, Lewis, Simpson, and Taggart prove that it remains NP-complete for diameter-two graphs, while Lewis, Cusack, and Dion prove the same for planar graphs (graphs that can be drawn in the plane without crossing edges) [17]. Remarkably, however, they also show that the problem is in P for diameter-two planar graphs! This line of research begs for more work.




Figure 10: A maximum r -path partition $\mathcal{P} = \{P^1, P^2, P^3, P^4, P^5\}$ of a tree T , with corresponding t -fold r -unsolvable configuration C_T shown.

For Question 3, Milans and Clark [18] were able to show that the decision problem “Is $\pi(G) \leq k$?” is Π_2^P -complete, a classification that is part of the polynomial-time hierarchy beyond NP. This makes graph

pebbling an excellent example of a problem that is beyond NP but short of EXPTIME-complete. Along these lines, Mendes, Pulaj, Wiedenbeck, and Yerger recently tied graph pebbling via bi-level integer programming to two-person Stackelberg games in economics.

Nonetheless, we still see results for various classes of graphs that the pebbling number can be calculated in polynomial time. For example, Herscovici, Hester, and Hurlbert proved that the pebbling number of a diameter two graph can be calculated in $O(n^4)$ time. But the first result of this kind is the formula for trees, which depends on constructing a decomposition of its edges into paths: a *path partition*. Suppose that $\mathcal{P} = \{P^1, \dots, P^m\}$ is a path partition of a tree T , with each P^i having length ℓ_i (i.e. $P^i \cong P_{\ell_i+1}$), written in non-increasing order. We say that \mathcal{P} is an *r-path partition* if r is an endpoint of P^1 and any other P^i that contains it, and that \mathcal{P} *majorizes* another *r-path partition* \mathcal{P}' if there is some j such that $\ell_i = \ell'_j$ for all $i < j$ and $\ell_j > \ell'_j$. If \mathcal{P} majorizes every *r-path partition* then it is *maximum* (see Figure 10). Chung proved the following theorem.

Theorem 7 (Chung [6]). *If ℓ_1, \dots, ℓ_m are the path lengths of a maximum *r-path partition* of a tree T then $\pi_t(T, r) = t2^{\ell_1} + \sum_{i=2}^m 2^{\ell_i} - m + 1$.*

One can imagine that the largest t -fold *r*-unsolvable configuration places all its pebbles on the non-*r* leaves of T ; indeed, this is so. Define the configuration C_T by placing $t2^{\ell_1} - 1$ pebbles on the leaf of P^1 , and $2^{\ell_i} - 1$ pebbles on the leaf of each remaining P^i . Then no path P^i ($i > 1$) has enough pebbles to reach its other endpoint, and P^1 cannot place t pebbles on r by itself. Chung’s proof used induction on n and t , involving the subtrees formed by components of $T - r$; other proofs used alternative induction steps (Moews), total weight [4], or the Weight Function Lemma [14], and also generalized the formula to include individual pebbling costs on each edge (Curtis-Hines-Hurlbert-Moyer). It is in Bunde et al. [4] where we find a linear algorithm to construct a maximum path partition for any tree.

This has led to results on other families of *chordal graphs*, defined by having no induced cycle of length more than 3, and are characterized by being a single vertex or having a *simplicial* vertex — a vertex whose neighbors form a complete subgraph — whose removal leaves a chordal graph. (A tree having no cycles, is trivially chordal.) Examples of graphs in this class whose pebbling numbers can be calculated in polynomial time include split graphs, 2-paths, and semi-2-trees (Alcón-Gutierrez-Hurlbert), the strong product of a path and a complete graph (Sieben), and graph powers of paths [1]. In some cases the graphs in a class may have more than one pebbling number formula to choose from. The case of split graphs involves six different formulas, depending on its structure. The algorithm runs in $O(n^\beta)$ steps, where $\beta = 2w/(w + 1) \cong 1.41$ and $w \cong 2.376$ is the exponent of matrix multiplication. A commonality, outside of split graphs, is that these graphs contain no induced pyramid.

Conjecture 8 (Alcón et al. [1]). *The pebbling numbers of pyramid-free chordal graphs can be calculated in polynomial time.*

Interestingly, Bender, Richmond, and Wormald proved that almost all chordal graphs are split graphs. Hence the pebbling numbers of almost all chordal graphs can be calculated in polynomial time.

For Question 2, there is very little work done so far. To state the question more precisely, suppose that G is a graph and C is a configuration of size $\pi(G, r)$ for some target vertex r . By definition, we know that C is *r*-solvable. Is it always possible, then, to compute an *r*-solution in polynomial time?

For trees the question is trivially yes. For any set of pebbling steps σ one can define the *transition digraph* G_σ , the multigraph of directed edges representing each pebbling step in σ . The No-Cycle Lemma of Moews states that if σ is an *r*-solution then there is some *r*-solution σ' for which $G_{\sigma'}$ is acyclic. This is a natural piece of intuition: a pebble traversing a cycle returns to its position, causing the loss of other pebbles for no gain. Thus, we may restrict our attention to *minimal r*-solutions: those for which the removal of any of its steps is no longer an *r*-solution. In the case of trees, then, minimal *r*-solutions only take steps toward r , decreasing the distance to r at each step — such solutions are called *greedy*.

Greedy solutions play an important role in calculating some pebbling numbers. For example, the *cost* of a solution is the number of pebbles used (equal to one more than the number of its pebbling steps). It can be shown that a minimal greedy solution costs at most 2^d , where d is the diameter of the graph — such solutions are called *cheap*. For example, a *2-path* is either a complete graph on at most three vertices or a chordal graph with exactly two simplicial vertices such that every maximal complete subgraph is a triangle. As is shown by Alcón et al., if G is a *2-path* on n vertices, with diameter d , then $\pi_t(G) = t2^d + n - d + 1$.

Part of their method shows that a configuration of size at least $\pi_2(G)$ has a greedy, hence cheap solution. Thus, $(t-1)2^d + n - d + 1$ pebbles remain to produce $t-1$ additional solutions by induction. There are numerous opportunities to apply this technique.

Unfortunately, not all solutions in chordal graphs are greedy (until the size of a configuration greatly exceeds the graph's pebbling number). However, a pebbling step is called *semi-greedy* if it moves a pebble from one vertex to a neighbor that is not of greater distance to r . Alc3n et al. [1] prove that all chordal graphs are *semi-greedy*: every configuration of size $\pi(G, r)$ on a chordal graph G has a semi-greedy r -solution. In addition, they prove another theorem reminiscent of the situation on trees. A *potential* vertex v is one containing at least two pebbles of a configuration. They prove that if G is chordal and C is an r -unsolvable configuration of maximum size, then every potential vertex is simplicial.

Chung's proof that cubes are Class 0 implicitly also proves that cubes are greedy. Her proof can easily be converted to an explicit algorithm that makes only greedy steps and that runs in $n \lg n$ time. This has direct applications to finding zero-sums in abelian groups in polynomial time. Algorithms for r -solving configurations of size $\pi(G)$ on graphs G in other graph classes would be useful.

4 Pebbling Variations

4.1 Target Pebbling: Satisfying Multiple Demands

Placing t pebbles on the same target vertex is just one level of generalization. We may instead be interested, for some *target distribution*² D , in placing $D(v)$ on each vertex v — i.e. *solving* D — from any configuration on G of sufficient size. From a supply-demand perspective, this may be the more natural optimization problem. The *target pebbling number* $\pi(G, D)$ is the minimum number of pebbles so that every configuration of this size is D -solvable. The first instance of this problem appeared in Crull et al. [8], in which $D(v) = 1$ for all v ; in this guise $\pi(G, D)$ was known as the *cover pebbling number*. A *stacked* configuration (or target distribution) is one that places all its pebbles (or targets) on a single vertex. A *positive* distribution D has $D(v) > 0$ for every v . Sj3strand generalized and proved the conjecture of [8] that, for any graph G , if D is a positive target distribution and C is a D -unsolvable configuration of maximum size then C is stacked. From this it is easy to calculate $\pi(G, D) = \max_v \sum_u 2^{\text{dist}(u,v)}$. Thus, the interesting case is when D is not positive.

Herscovici et al. conjecture for every graph G that $\pi(G, D) \leq \pi_t(G)$ for every D of size t . The intuition is that, if $|C| \geq \pi_t(G)$ then C can place t pebbles on any vertex of G — it seems reasonable, then, to believe that C can place t pebbles on any combination of t vertices. It feels a little bit like a convexity notion. The authors verify the conjecture for any G with $\pi(G) = 2^{\text{diam}(G)}$, such as complete graphs, even cycles, and cubes, as well as for trees and odd cycles. This has come to be known as the *Weak Target Conjecture*, and was also verified recently by Hurlbert and Seddiq for 2-paths and Kneser graphs $K(m, 2)$.

Upon closer inspection of a graph like the complete graph, for example, every maximum size t -fold r -unsolvable configuration has potential $t-1$ and thus an odd number of pebbles on each vertex except r ; in particular, $C(r) = 0$ and $C(v) > 0$ for every $v \neq r$. For general D , such extremal configurations must have $C(u) = 0$ for every $u \in D$, and so each extra target in D drops $\pi(G, D)$ by one. As with configurations, define the *support* of a distribution, $\text{supp}(D)$, to be the set of vertices u with $D(u) > 0$, with $s(D) = |\text{supp}(D)|$. Then the following conjecture, true for complete graphs, seems natural.

Conjecture 9 (Strong Target Conjecture [1]). *For every graph G and target distribution D of size t we have $\pi(G, D) \leq \pi_t(G) - s(D) + 1$.*

The authors of [1] verify this conjecture for trees and for powers of paths. It would be interesting to explore trees further: for example, one could imagine that an alternative partitioning of a tree T might yield an explicit formula for $p(T, D)$. Even the case $|D| = s(D) = 2$ is unknown.

Finally, one also finds in [1] another motivation for studying target distributions that relates to Graham's conjecture. There they posit the generalization that $\pi(G_1 \square G_2, D_1 \times D_2) \leq \pi(G_1, D_1)\pi(G_2, D_2)$, which also implies conjectures of Herscovici and of Lourdasamy. As is often the case, finding just the right generalization of a problem can be the key to cracking it.

²Configurations and distributions are both functions $V(G) \rightarrow \mathbb{N}$. However, to avoid confusion, we use configurations for pebbles and distributions for targets.

4.2 Optimal Pebbling: Finding the Best Configuration

Here we change the paradigm somewhat. So far we have been considering the worst-case scenario: the largest configuration that can't solve some target. Now we consider the best-case scenario: the smallest configuration that can solve every target. The former is like a two-person game in which your adversary places the pebbles; optimal pebbling is a game of solitaire in which you get to place your own pebbles. It's quite similar to finding the best placement of cops to capture any robber, or determining the smallest distance- d dominating set. Pachter et al. were the first to define the *optimal pebbling number* of a graph G , denoted $\pi^*(G)$, to be the smallest m for which there is some configuration C of size m that is r -solvable for every r . Deciding if $\pi^*(G) \leq k$ is easier than the same question for $\pi(G)$, but it is still NP-complete [18]. In this setting, however, Shiue showed that the cartesian product bound does hold:

$$\pi^*(G \square H) \leq \pi^*(G)\pi^*(H).$$

Pachter et al. showed that $\pi^*(P_{3k+r}) = 2k + r$ for all $k \geq 0$ and all $0 \leq r \leq 2$ (Friedman and Wyels showed that the same result holds for cycles). Basically, you place two pebbles on every third vertex on the path and add a pebble on any remaining vertex, and then show by induction that this is best possible. Bunde et al. [4] carefully generalized the construction to any tree T , yielding $\pi^*(T) \leq \lceil 2n/3 \rceil$, and thus $\pi^*(G) \leq \lceil 2n/3 \rceil$ for all G . Note that every vertex without a pebble in this path configuration C is adjacent to a vertex with a pebble — that is, $\text{supp}(C)$ is a *dominating set*. More generally, define $\gamma(G)$ to be the minimum size of a dominating set in G . Then we always have $\pi^*(G) \leq 2\gamma(G)$. Thus, finding dominating sets of various types has been the dominant (sorry, we couldn't resist that) method in this area, and Chellali et al. [5] connect variations of optimal pebbling to several different domination parameters.

In particular, a *distance k dominating set* S is a set of vertices of a graph G such that every vertex of G is within distance k from some vertex of S . The size of the smallest such set is denoted $\gamma_k(G)$, and this yields the more general bound

$$\pi^*(G) \leq 2^k \gamma_k(G). \tag{3}$$

An error-correcting code that corrects up to k errors is an example of a distance k dominating set in the cube. In fact, Moews used the result of Cohen that $\gamma(Q^d) \leq 2^m$, where $m \leq d(1 - H(k/d)) + 2 \lg d$ and $H(x) = -x \lg x - (1-x) \lg(1-x)$, to prove that $\pi^*(Q^d) \leq 2d^2(4/3)^d$. This has since been improved by Fu, Huang, and Shiue to

$$\pi^*(Q^d) = O\left(d^{3/2}(4/3)^d\right).$$

For the lower bound $\pi^*(Q^d) \geq (\frac{4}{3})^d$, Moews introduces a continuous optimal pebbling version, with real-valued configurations and real-valued pebbling steps, still with 50% cost, and with corresponding parameter $\hat{\pi}^*$. For example, $\hat{\pi}^*(Q^1) = 4/3$: place $2/3$ on each vertex — placing any other values a and b with smaller sum will fail to move 1 to the smaller of the two. Also, since integer configurations are real configurations, we see that $\hat{\pi}^*(G) \leq \pi^*(G)$ for all G . Moews then proves the corresponding stronger version of Graham's conjecture:

$$\hat{\pi}^*(G \square H) = \hat{\pi}^*(G)\hat{\pi}^*(H)$$

to obtain the result. Finally, he follows Cohen's probabilistic approach — showing that a randomly chosen set of the right size will be a distance k dominating set with positive probability — to prove that every graph G satisfies $\pi^*(G^d) = \hat{\pi}^*(G)^{d+O(\lg d)}$, where G^d denotes the Cartesian product of G with itself d times.

It is natural to study how $\pi^*(G)$ decreases as the minimum degree of G increases. Let $N_k(v)$ denote the set of vertices within distance k of v , having size $\text{deg}_k(v)$, and define $\delta_k(G) = \min_v \text{deg}_k(v)$. Write $N_k(U) = \cup_{u \in U} N_k(u)$. Typically, we avoid writing the subscript when $k = 1$. In [4] we find a clever argument of Czygrinow that every graph G has $\gamma_{2k}(G) \leq n/\delta_k(G)$. Indeed, construct a distance- $2k$ dominating set S starting with any vertex $v \in G$, and initialize $T = N_k(v)$. If T is not a distance- k dominating set then add any vertex $u \notin N_k(T)$ to S and add $N_k(u)$ to T (this is a disjoint union). By the time S has size $n/\text{deg}_k(G)$, T is a distance- k dominating set, making S a distance- $2k$ dominating set. For instance, he applies this bound to that of (3) to obtain

$$\pi^*(G) \leq 4\gamma_2(G) \leq 4 \left(\frac{n}{\delta + 1} \right).$$

Czygrinow, Hurlbert, Katona, and Papp show that the exact bound is not attainable, although it is sharp: for every $\epsilon > 0$ there is a diameter two graph G (specifically $\overline{K_{\sqrt{n}} \square K_{\sqrt{n}}}$) with $\pi^*(G) > (4 - \epsilon) \binom{n}{\delta+1}$. However, they prove that when $\text{diam}(G) \geq 3$ the coefficient in the upper bound can be improved to 3.75. They also display, for any $\epsilon > 0$, infinitely many graphs with coefficient at least $\frac{8}{3} - \epsilon$, leaving a small gap for readers to close.

One interesting problem that surprisingly remains unresolved is the following. Györi et al. [11] proved that

$$(2/13)mn \leq \pi^*(P_m \square P_n) \leq (2/7)mn,$$

but the correct bound is unknown. A final pursuit that is unexplored to date is investigating $\pi^*(G, D)$ for some target D — that is, maybe you don't need to potentially reach every vertex, and maybe you might need to reach some vertices multiple times.

4.3 Threshold Pebbling: The Random Case

Now that we've discussed the worst- and best-case scenarios, let's think about the "typical" case, as initiated by Czygrinow, Eaton, Hurlbert, and Kayll.

In this paradigm, neither the pebbler nor their adversary controls the initial configuration C ; instead it is chosen at random. Naturally, large configurations will be solvable with high probability, while small configurations will be solvable with low probability. One might wonder if there is some threshold, i.e., phase transition, at which the transition from low to high probability is sharp — indeed, there is.

Before we dive into the details, let us informally illustrate that threshold pebbling is a generalization of an unlabeled version of Feller's well-known Birthday Problem! Consider the complete graph K_n . We want to know: "If pebbles are randomly placed on the vertices, how many pebbles are needed (as a function of n) so that the probability of pebbling to any possible vertex tends to 1?". Indeed, a configuration of size less than n on K_n is solvable if and only if some vertex has at least two pebbles. Just as with the birthday problem, this requires $\Theta(\sqrt{n})$ pebbles, and so a pebbling threshold function for the family of complete graphs is $\Theta(\sqrt{n})$.

To make this formal, consider an infinite sequence of graphs $\mathcal{G} = (G_1, G_2, \dots)$ with an increasing number of vertices, and a sequence of positive integers (m_1, m_2, \dots) . Define $\mathbf{p}(G_i, m_i)$ to be the probability that C , a randomly chosen configuration of m_i pebbles is *solvable* on G_i (i.e. r -solvable for all r). Note that, as opposed to randomly placing each pebble independently, according to some fixed probability, we select C uniformly at random from the set of all configurations on $V(G)$ of a fixed size m . A *pebbling threshold function* of \mathcal{G} , $t(k)$, is a function on \mathbb{N} such that, as $k \rightarrow \infty$, $\mathbf{p}(G_k, m_k) \rightarrow 0$ (resp. 1) whenever $m_k/t(k) \rightarrow \infty$ (resp. 0). We denote by $\tau(\mathcal{G})$ the set of all thresholds for \mathcal{G} .

Of course, there is no reason, *a priori*, that a pebbling threshold even exists for any particular sequence of graphs! However, we have the following theorem.

Theorem 10 (Bekmetjev et al. [3]). *Every infinite sequence of graphs has a pebbling threshold function.*

The proof of this theorem utilizes beautiful results from extremal set theory, so we sketch it below. For a graph G , define the family \mathcal{F} of all unsolvable configurations on G , with \mathcal{F}_m being those of size m . We may view each configuration as a multiset of m vertices, of which there are $\binom{n}{m} = n(n+1) \cdots (n+m-1)/m!$. Then, we see that $\mathbf{p}(C) = |\mathcal{F}_m|/\binom{n}{m}$ when $|C| = m$. For any family \mathcal{A} of multisets of $V(G)$, define the *shadow* $\partial\mathcal{A}_m = \{A - v \mid v \in A \in \mathcal{A}_m\}$. Note that $\partial\mathcal{F}_m \subseteq \mathcal{F}_{m-1}$, which is what we will mean by *monotone*. In the lattice of subsets, the famed Kruskal-Katona Theorem gives a precise formula for the minimum shadow size of a family \mathcal{B}_m of m -subsets of size M , based on something called the "cascade representation" of $|\mathcal{B}_m|$, which we won't get into. The formula is beautiful but difficult to use in calculations. For this, we rely on the excellent approximation of Lovász, which states for any real x that if $|\mathcal{B}_m| = \binom{x}{m} = x(x-1) \cdots (x-m+1)/m!$ then $|\partial\mathcal{B}_m| \geq \binom{x}{m-1}$. Then Bollobás and Thomason used this to show for monotone \mathcal{B} that if $j < k$ then $\mathbf{p}(\mathcal{B}_j)^k \geq \mathbf{p}(\mathcal{B}_k)^j$, where $\mathbf{p}(\mathcal{B}_m) = |\mathcal{B}_m|/\binom{n}{m}$. From this it follows that, if we apply these calculations to an infinite sequence of graphs (G_1, G_2, \dots) , with increasing $n(G_i)$ and families $\mathcal{B}_m(i)$ defined on G_i with m a function of n , then there is some threshold $t = t(\mathcal{B})$ for which, as $i \rightarrow \infty$, $\mathbf{p}(\mathcal{B}_m) \rightarrow 1$ when $m/t \rightarrow 0$ and $\mathbf{p}(\mathcal{B}_m) \rightarrow 0$ when $m/t \rightarrow 1$.

The Kruskal-Katona theorem was successfully generalized to the (infinite) lattice of multisets by Clements and Linström. Then Bekmetjev et al. [3] proved the analogue of Lovász's inequality — if \mathcal{A} is a family of

multisets with $|\mathcal{A}_m| = \binom{x}{m}$ then $|\partial\mathcal{A}_m| \geq \binom{x}{m-1}$ — and used this to settle the existence of a threshold $\tau(\mathcal{A})$ for any monotone family \mathcal{A} of multisets. Theorem 10 states this in particular for the monotone families of unsolvable configurations of a sequence of graphs.

Regarding specific examples of graph sequences, recall that $\tau(K_n) = \Theta(\sqrt{n})$. At the other end of the spectrum is the sequence of paths \mathcal{P} . A series of papers resulted in the lower (Czygrinow-Hurlbert) and upper [3] bounds $\tau(\mathcal{P}) \subseteq \Omega(n2^{c\sqrt{\lg n}}) \cap O(n2^{C\sqrt{\lg n}})$ for any $c < 1 < C$, and finally Moews recently found the answer:

$$\tau(\mathcal{P}) = \Theta\left(n2^{\sqrt{\lg n}}/\sqrt{\lg n}\right).$$

It is not difficult to see that every graph threshold sits between $\tau(\mathcal{K})$ and $\tau(\mathcal{P})$, and Moews proved that, for every function g in this range, there is a graph sequence \mathcal{G} with $\tau(\mathcal{G}) = \Theta(g)$. Czygrinow, Eaton, Hurlbert, and Kayll showed that if \mathcal{G} is a sequence of graphs having bounded diameter, though, then $\tau(\mathcal{G}) \in O(n)$, while if the graphs have minimum degree δ , for $\sqrt{n} \ll \delta < n$, then $\tau(\mathcal{G}) \in O(n^{3/2}/\delta)$ (which yields $\Theta(\sqrt{n})$ for $\delta \in \Theta(n)$) (Czygrinow-Hurlbert).

Interestingly, Björklund and Holmgren proved that π and τ are not always positively correlated. A *blow-up* of a vertex v in a graph replaces v by a complete graph, each vertex of which is adjacent to each vertex in $N(v)$. For example, blowing up an endpoint of an appropriate length path will realize the required graph sequence in Moews' density result above. Here, Björklund and Holmgren define G_n to be the n -vertex graph blown up from an endpoint of a length $\lg n$ path, H_n to be the n -vertex graph blown up from a midpoint of a length $\lg n + c$ path, for some large constant c , and show that $\pi(G_n) \leq 3n \ll \pi(H_n)$. Then, for $\mathcal{G} = (G_n)_{n \geq 1}$ and $\mathcal{H} = (H_n)_{n \geq 1}$, they prove that $\tau(\mathcal{H}) \in O(n^{0.8})$, while $\tau(\mathcal{G}) \in \Omega(n^{0.99})$. They also give more complicated and extreme examples of \mathcal{G} and \mathcal{H} with $\pi(G_n) \ll \pi(H_n)$, $\tau(\mathcal{H})$ arbitrarily close to $\tau(\mathcal{K})$, and $\tau(\mathcal{G})$ arbitrarily close to $\tau(\mathcal{P})$.

Threshold arguments are typically local. For upper bounds one usually shows that a random configuration of many pebbles almost always places sufficiently many pebbles in every small subgraph. For lower bounds one typically shows that a random configuration of few pebbles almost always contains a large hole — some $N_k[v]$ with no pebbles in it — and cannot move enough pebbles to the boundary of the hole sufficient to reach v . The details of what small and large means and what probabilistic techniques are used vary upon the instance.

Along the lines of Cartesian products, both Alon and Czygrinow-Wagner produced bounds for the sequence $\mathcal{Q} = (Q^d)_{d \geq 1}$ of cubes:

$$\tau(\mathcal{Q}) \in \Omega(n^{1-\epsilon}) \cap O(n/(\lg \lg n)^{1-\epsilon}) \text{ for all } \epsilon > 0.$$

This has room for tightening. For the sequence of products of complete graphs $\mathcal{K}^2 = (K_k^2)_{k \geq 1}$, Bekmetjev and Hurlbert proved that $\tau(\mathcal{K}^2) = \Theta(\sqrt{n})$. This satisfies the following natural threshold version of Graham's conjecture, which needs to be stated carefully to account for the re-scaling of n (i.e. $n(G_k \square H_k) = n(G_k)n(H_k)$).

Conjecture 11 (Hurlbert [12]). *Let $\mathcal{F} = (F_k)_{k \geq 1}$, $\mathcal{G} = (G_k)_{k \geq 1}$, and $\mathcal{H} = (H_k)_{k \geq 1}$, where $F_k = G_k \square H_k$, and suppose that $f \in \tau(\mathcal{F})$, $g \in \tau(\mathcal{G})$, and $h \in \tau(\mathcal{H})$. Then*

$$f(n(F_k)) \in O(g(n(G_k))h(n(H_k))).$$

Two additional directions are worthy of exploration. One is to change the model of generating a random configuration. Unlike in the random graph case, in which the hitting time, static, and Erdős-Reyni models all concur, a different configuration model can change the threshold. For example, Czygrinow and Hurlbert showed that if the pebbles are labeled and each pebble gets to choose its own vertex, then the path threshold drops to $\Theta(n \lg n / \lg \lg n)$. Finally, one can also consider the threshold for cover pebbling, as is done by Godbole, Watson, and Yeger, where they prove the following sharp result. Let $\phi = (1 + \sqrt{5})/2$ and $\omega = \omega(n) \rightarrow \infty$. Then a uniformly random configuration on K_n is almost surely cover-solvable when $|C| \geq \phi n + \omega\sqrt{n}$ and almost surely not cover-solvable when $|C| \leq \phi n - \omega\sqrt{n}$.

5 Final Comments

The challenges involved in graph pebbling touch different areas of mathematics beyond graph theory including number theory, optimization, computation, extremal set theory, and probability. The Lemke-Kleitman

Conjecture 3 remains open in the general case. Pebbling has also played a role in similar, combinatorial number-theoretic problems as well. For example, Knapp [16] recently used a pebbling technique to show the existence of 2-adic solutions to certain homogeneous additive equations! From a different perspective, the computational challenges are daunting; any progress toward efficiently computing such parameters would enable great progress in graph pebbling. These are, of course, in addition to Graham’s Conjecture.

Regardless of your background, we invite you and your students to pebble! Clever applications and connections to other areas of mathematics await!

References

- [1] ALCÓN, L., AND HURLBERT, G. Pebbling in powers of paths. *in preparation* (2021).
- [2] ASPLUND, J., HURLBERT, G., AND KENTER, F. Pebbling on graph products and other binary graph constructions. *Australas. J. Combin.* 71, 2 (2018), 246–260.
- [3] BEKMETJEV, A., BRIGHTWELL, G., CZYGRINOW, A., AND HURLBERT, G. Thresholds for families of multisets, with an application to graph pebbling. *Discrete Math.* 269, 1–3 (2003), 21–34.
- [4] BUNDE, D. P., CHAMBERS, E. W., CRANSTON, D., MILANS, K., AND WEST, D. B. Pebbling and optimal pebbling in graphs. *J. Graph Theory* 57, 3 (2008), 215–238.
- [5] CHELLALI, M., HAYNES, T. W., HEDETNIEMI, S. T., AND LEWIS, T. M. Restricted optimal pebbling and domination in graphs. *Discrete Appl. Math.* 221 (2017), 46–53.
- [6] CHUNG, F. Pebbling in hypercubes. *SIAM J. Discrete Math.* 2, 4 (1989), 467–472.
- [7] CRANSTON, D. W., POSTLE, L., XUE, C., AND YERGER, C. Modified linear programming and Class 0 bounds for graph pebbling. *J. Comb. Optim.* 34, 1 (2017), 114–132.
- [8] CRULL, B., CUNDIFF, T., FELTMAN, P., HURLBERT, G. H., PUDWELL, L., SZANISZLO, Z., AND TUZA, Z. The cover pebbling number of graphs. *Discrete Math.* 296, 1 (2005), 15–23.
- [9] CUSACK, C. A., POWERS, M., AND BEKMETJEV, A. Doppelgangers and Lemke graphs. *Discrete Math.* 341, 10 (2018), 2686–2693.
- [10] CZYGRINOW, A., HURLBERT, G., KIERSTEAD, H., AND TROTTER, W. T. A note on graph pebbling. *Graphs Combin.* 18 (2002), 219–225.
- [11] GYÓRI, E., KATONA, G. Y., AND PAPP, L. Optimal pebbling number of the square grid. *Graphs Combin.* 36, 3 (2020), 803–829.
- [12] HURLBERT, G. General graph pebbling. *Discrete Appl. Math.* 161, 9 (2013), 1221–1231.
- [13] HURLBERT, G. Graph Pebbling. *Handbook of Graph Theory (2nd ed.)*, Ed: J.L. Gross, J. Yellen, P. Zhang, Chapman and Hall/CRC, Kalamazoo (2014), 1428–1449.
- [14] HURLBERT, G. The weight function lemma for graph pebbling. *J. Combin. Opt.* 34, 2 (2017), 343–361.
- [15] KENTER, F., SKIPPER, D., AND WILSON, D. Computing bounds on product graph pebbling numbers. *Theoret. Comput. Sci.* 803 (2020), 160–177.
- [16] KNAPP, M. P. Distance pebbling on directed cycle graphs. *J. Combin. Math. Combin. Comput.* 111 (2019), 283–286.
- [17] LEWIS, T., CUSACK, C., AND DION, L. The complexity of pebbling reachability and solvability in planar and outerplanar graphs. *Discrete Appl. Math.* 172 (2014), 62–74.
- [18] MILANS, K., AND CLARK, B. The complexity of graph pebbling. *SIAM J. Discrete Math.* 20, 3 (2006), 769–798.

- [19] MOEWS, D. Pebbling graphs. *J. Combin. Theory Ser. B* 55, 2 (1992), 244–252.
- [20] POSTLE, L., STREIB, N., AND YERGER, C. Pebbling graphs of diameter three and four. *J. Graph Theory* (2012).