

An Exploration of Graph Pebbling

This is the master thesis of *Herman Bergwerf*, who was in the *Mathematical Foundations of Computer Science* programme, and whose student number is *1047633*. This thesis was supervised by *Wieb Bosma*, at the Faculty of Science.

Preface

I wrote this thesis for my master's degree at Radboud University. It is the culmination of a five-year journey into mathematics and formal logic, which was initiated by a side project where I attempted to build my own theorem prover. Here I want to briefly tell a story about this journey.



When I first went to university eight years ago I had little interest in mathematics. I wanted to understand more about the machinery of life and decided to join the *Nanobiology* program in Delft, which offered a glimpse into the complex world of low-level biological systems. During this time I worked on a side project called *QEDb*, which aimed to develop a structured model and an attractive interface for writing and storing mathematical derivations. Initially I had solutions to physics problems in mind, such as the exercises we had to do for our classes. The derivation database would be a way for students to share their solutions in a rigorous way, and to learn from the solutions of others. I worked on the *QEDb* project for several years, but, as I learned more about mathematical proofs, I kept finding things that required fundamental changes to my program. At first it only reasoned with equality and rewriting, then I implemented deduction rules, and at some point I learned about quantifiers and realized they had to be expressible in the system too. Eventually I abandoned the project, but it had sparked a definite interest in mathematics and formal systems.

I had known about the Coq theorem prover for some time, but I didn't have a clue about how it worked or how I could learn it. I thought that the most effective way to achieve my goal was to build my own formal system. A course from Robbert Krebbers, where I first learned how to use Coq, changed this. I realized that Coq solved many problems I had encountered, and that there was a great deal for me to learn about theorem provers. This was an important motivation to study theoretical computer science in Nijmegen, where courses were offered about type systems, proof assistants, formal semantics, and other interesting topics. I was also increasingly fascinated by purely mathematical problems. Some theorems, like Gödel's incompleteness theorem, or the decidability of linear integer arithmetic, were surrounded by an air of mystery, and I *had* to know more about these things; I wanted to know how such results were even *possible*. I particularly enjoyed the courses of Wim Veldman, whose classes were full of interesting stories and cliffhangers. His course about the independence of the continuum hypothesis from ZFC set theory is perhaps the most sophisticated and intriguing bit of theory I have ever studied.

For my thesis I wanted to dive into a similarly mystifying topic, and do something that is not directly related to computer programming. When I asked Wim Veldman how to find an interesting problem to work on, he told me to read about a subject I liked and just ask myself questions about it, no matter how simple. He recommended the book *The Strange Logic of Random Graphs* by Joel Spencer, which among other things treats Ehrenfeucht-Fraïssé games. In a related article I encountered the term *pebbling game* that, quite by accident, brought me to the topic of this thesis (which is actually completely unrelated to what that article was talking about). For over half a year I puzzled feverishly over various problems related to graph pebbling, eager to find a gold nugget that would reveal the path to a new solution. In the end I wasn't able to solve any of the open problems, but I got the opportunity to develop my own intuition, and address some small questions along the way. In this thesis, I want to show what a rich and fascinating topic graph pebbling is. I hope it will inspire others to study it too!



There is a number people that played a role in this process and who I am especially grateful to. First of all I want to thank Wieb Bosma, who despite being unfamiliar with graph pebbling was prepared to supervise my thesis. We spent a lot of time discussing various results, and Wieb was always open to exploring new directions. He gave me a lot of detailed feedback that sharpened the style and quality of my writing. Without him I would not have reached this point. I am also very thankful to Wim Veldman, who through his classes inspired a certain way of thinking and writing about mathematics. His enthusiasm and rich use of metaphors were very motivating, and I never so enjoyed sitting through lectures. His course about intuitionistic mathematics revealed a whole new way of looking at logic. Finally, I also want to thank Robbert Krebbers, from whom I learned developing efficient Coq proofs, and who has, over the past years, answered many of my emails about Coq.

Herman Bergwerf
April 2023, Nijmegen

Contents

1	Introduction	7
2	Pebbling Numbers	9
3	Weight Functions	19
4	Pebble Flows	27
5	Hypercubes	33
6	Zero-sums	43
	Bibliography	51
	Appendix	53

Notation

- The set of *natural numbers*, including the number zero, is denoted by \mathbb{N} .
- The set of *non-negative rational numbers* is denoted by \mathbb{Q}^+ .
- The *number of elements* of a set S is written as $\#S$.
- The *domain* of a function f is written as $\text{dom}(f)$.
- The *support* of a real-valued function f is defined as $\text{supp}(f) := \{x \in \text{dom}(f) \mid f(x) \neq 0\}$.
- The *size* of a real-valued function f is defined as $|f| := \sum_{x \in \text{dom}(f)} f(x)$.
- The *element-wise addition* of two real-valued functions f and g is written as $f + g$.
- The *element-wise subtraction* of two real-valued functions f and g is written as $f - g$.
- The *element-wise multiplication* of two real-valued functions f and g is written as $f \cdot g$.
- The *complete graph* on n vertices is denoted by K_n .
- The *cycle graph* on n vertices is denoted by C_n .
- The *path graph* on n vertices is denoted by P_n .

1 Introduction

The topic of this treatise is a combinatorial technique called *Graph Pebbling*. Before we start investigating this, it is worth considering where such a thing comes from. In the early 20th century one of the great unsolved problems in mathematics was the *Entscheidungsproblem*, which asked if there exists a procedure to determine if a given sentence in first-order logic is universally true. One of the people who worked on this was Frank Ramseyⁱ, who found such a procedure for a restricted class of logical formulaeⁱⁱ. It is now well known that a general decision procedure for first-order logic does *not exist*, which was shown by both Churchⁱⁱⁱ and Turing^{iv} using formal systems to express *algorithms*. Ramsey didn't witness this result; he passed away at just 26. In his work on the Entscheidungsproblem he introduced a curious theorem about graphs that carried an intriguing concept; certain orderly substructures are present in any sufficiently large structure, *no matter how disordered*. This result inspired a whole mathematical field that is now known as *Ramsey Theory*.

The renowned mathematician Paul Erdős helped to popularize Ramsey Theory, and worked on many problems in this new field. In 1961 he showed^v, together with Ginzburg and Ziv, that every integer sequence of length $2n - 1$ has a subsequence of n integers that together sum to a multiple of n , *i.e.* zero modulo n . This seminal work inspired research into other *zero-sum* problems. Communication between Erdős and the mathematician Paul Lemke resulted in the theorem stated below, which we shall call the *Erdős-Lemke conjecture*.

Theorem 1.1. (Erdős-Lemke conjecture) *Given positive integers n, d , and a_1, a_2, \dots, a_d with $d \mid n$ and $a_i \mid n$ for $i = 1, 2, \dots, d$, there is a non-empty subset S of $\{1, 2, \dots, d\}$ such that:*

$$\sum_{i \in S} a_i \equiv 0 \pmod{d} \quad \text{and} \quad \sum_{i \in S} a_i \leq n.$$

A proof of this conjecture was published by Lemke and Kleitman in [2] a few years after its conception. Their proof, which uses induction on the number of prime factors of d , is quite complicated and difficult to divide into smaller steps. Fan Chung developed an alternative proof in [3] based on a game of moving pebbles between the vertices of a graph, an idea originating from Lagarias and Saks (who are absent in the literature). In this proof the various reasoning steps are more clearly separated, and most work is spent analyzing the pebbling game on a hypercube. Chung raised some interesting questions about the pebbling game on other graphs, inspiring a string of further research into what is now called Graph Pebbling^{vi}.

ⁱFor a brief outline of Ramsey's life, see [1].

ⁱⁱ"On a Problem of Formal Logic" (1930)

ⁱⁱⁱ"An Unsolvability Problem of Elementary Number Theory" (1936)

^{iv}"On Computable Numbers, with an Application to the Entscheidungsproblem" (1936)

^v"Theorem in the Additive Number Theory" (1961)

^{vi}For a recent survey of graph pebbling research, see [20].




Figure 1 – *Man with Cuboid* by M.C. Escher (1958)

Let's consider an instance of this pebbling game. Suppose I distribute seven pebbles over the vertices of a cube, where it is possible to put multiple pebbles on a single vertex. Starting from this initial configuration of pebbles, you are allowed to make changes with the following rule: You may add *one* pebble to any vertex in exchange for removing *two* pebbles from an adjacent vertex, resulting in a net loss of one pebble. This is called a *pebbling step*. Your objective is to put one pebble on a designated empty vertex using pebbling steps. Can I pick a certain initial configuration with an empty vertex for which you are unable to achieve the objective? And what happens if I distribute eight pebbles instead of seven?

We are going to explore this game in detail in the coming chapters. In the next chapter we will first define some graph pebbling concepts formally, and get more familiar with them by looking at pebbling in trees and diameter-2 graphs. We will also discuss some questions that reveal the complexity of our topic. In Chapter 3 we explore a technique that uses linear inequalities to determine if an initial pebble configuration can be solved by a specific strategy, yielding a number of elegant proofs and a method to apply linear programming. In Chapter 4 we think about the direction in which pebbles need to be moved, and we prove that back-and-forth movement between vertices is never a necessity. In Chapter 5 we review Chung's results about pebbling on a hypercube, and we describe a more detailed proof with a more general induction argument. Finally, in Chapter 6, we will prove the Erdős-Lemke conjecture via graph pebbling. Our investigation of this result was driven by a desire to describe a formal proof using the *Coq Proof Assistant*. Details of the completed formalization are described in the appendix.

2 Pebbling Numbers

Contrary to most of the literature about graph pebbling, we are going to describe the pebbling game using directed graphs and edge weights. This will be particularly useful for our proof of the Erdős-Lemke conjecture.

Definition 2.1. A simple graph G is a pair (V, E) of a set of vertices V and a set of edges E , where each edge $\{u, v\} \in E$ is a set of two distinct vertices. The vertices of G are written as $V(G)$ and the edges as $E(G)$.

Definition 2.2. An edge-weighted digraph G is a triplet (V, E, ω) of a set of vertices V , a set of edges $E \subseteq V \times V$, and an edge-weight function $\omega : E \rightarrow \mathbb{N}$ that assigns a weight to every edge. The vertices of G are written as $V(G)$, the edges as $E(G)$, and the edge-weight function as ω_G . An edge $(u, v) \in E$ begins at u and ends at v .

We refer to edge-weighted digraphs simply as *graphs*. Sometimes graphs with undirected edges are given, in which case you may assume that all edges are bidirectional. We define a special notation to convert a simple graph, such as the complete graph K_n , to an edge-weighted digraph where all edges are bidirectional, and the weight of every edge is k :

Definition 2.3. Let G be a simple graph and let $k \in \mathbb{N}$. The graph $G^{(k)}$ is defined as (V, E, ω) where $V = V(G)$, $E = \{(u, v) \mid \{u, v\} \in E(G)\}$, and $\omega(u, v) = k$ for all $u, v \in V(G)$.

Throughout this text, we will encounter a special way to construct larger graphs from smaller graphs known as a *Cartesian graph product*. The Cartesian product of two graphs G and H is written as $G \square H$. The box symbol is a reference to the orthogonality of the construction; the product of two paths is a grid, and the product of two cycles is a torus.

Definition 2.4. Let G and H be graphs. The product $G \square H$ is defined as (V, E, ω) where:

$$V = V(G) \times V(H)$$

$$E = \left\{ \begin{array}{l} ((u_1, u_2), (v_1, v_2)) \\ \in V \times V \end{array} \left| \begin{array}{l} ((u_1, v_1) \in E(G) \wedge u_2 = v_2) \vee \\ ((u_2, v_2) \in E(H) \wedge u_1 = v_1) \end{array} \right. \right\}$$

$$\omega((u_1, u_2), (v_1, v_2)) = \begin{cases} \omega_G(u_1, v_1) & \text{if } (u_1, v_1) \in E(G) \\ \omega_H(u_2, v_2) & \text{if } (u_2, v_2) \in E(H) \end{cases}$$




Figure 2 – The product of two cycles is a torus.

Graph pebbling starts with a distribution of pebbles over the vertices of a graph, called the *initial configuration*. This configuration can be altered with *pebbling steps*: If an edge (u, v) has weight k , and u has at least k pebbles, then we may remove k pebbles from u and add one pebble to v . Essentially all of the literature about graph pebbling defines pebbling steps to require a standard amount of two pebbles, i.e. $k = 2$ for all edges. We will in some cases refer to this as *weight-2 pebbling*. To avoid trivialities, we assume that edges always have a weight of at least two; $k \geq 2$.

Definition 2.5. Let G be a graph. A configuration c on G is a function $V(G) \rightarrow \mathbb{N}$.

Definition 2.6. Let G be a graph. The pebbling step relation $c_1 \rightarrow c_2$ between configurations c_1 and c_2 on G holds if there is an edge (u, v) such that $c_1(u) \geq k$ for $k = \omega_G(u, v)$, and c_2 results from c_1 after k pebbles are removed from u and one pebble is added to v :

$$\forall w \in V(G) : c_2(w) = \begin{cases} c_1(u) - k & \text{if } w = u \\ c_1(v) + 1 & \text{if } w = v \\ c_1(w) & \text{otherwise} \end{cases}$$

We use an asterisk to denote the reflexive-transitive closure of this relation: $c_1 \rightarrow^* c_n$ means that there is a sequence of zero or more pebbling steps $c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_n$ from configuration c_1 to c_n . If, starting from a configuration c , it is possible to put n pebbles on a target vertex t via pebbling steps, then we say that c is n -fold t -solvable, or that n pebbles can be *moved*ⁱ to t . When $n = 1$ we just write “ t -solvable”.

Definition 2.7. Let G be a graph. A configuration c on G is n -fold t -solvable if there is a configuration c^* such that $c \rightarrow^* c^*$ and $c^*(t) \geq n$. As a first-order formula:

$$\text{Solv}_{n,t}(c) := \exists c^* : c \rightarrow^* c^* \wedge c^*(t) \geq n$$

ⁱHere ‘moving’ n pebbles also describes cases where the initial configuration already puts some number of pebbles on t , such that less or no additional pebbles need to be added to reach n .

Definition 2.8. Let G be a graph and let $t \in V(G)$. The n -fold t -pebbling number $\pi_n(G, t)$ is the smallest number p such that every configuration c with $|c| \geq p$ is n -fold t -solvable. The pebbling number $\pi(G)$ of G is defined as the largest 1-fold pebbling number among its vertices.

$$\begin{aligned}\pi_n(G, t) &:= \min \{p \in \mathbb{N} \mid \forall c : |c| \geq p \Rightarrow \text{Solv}_{n,t}(c)\} \\ \pi(G, t) &:= \pi_1(G, t) \\ \pi(G) &:= \max \{\pi(G, v) \mid v \in V(G)\}\end{aligned}$$

Let's use these definitions to prove some simple theorems.

Theorem 2.9. For a simple graph G holds $\pi(G^{(2)}) \geq \#V(G)$.

Proof. Note that the weight of all edges in $G^{(2)}$ is 2 by definition. Pick any vertex t of G and put a single pebble on each vertex except t . The resulting configuration contains $\#V(G) - 1$ pebbles and is not t -solvable since no pebbling steps can be applied. \square

Theorem 2.10. For $n \geq 1$ and $k \geq 1$ holds $\pi(K_n^{(k)}) = (n - 1)(k - 1) + 1$.

Proof. Note that all vertices of $K_n^{(k)}$ have the same pebbling number. Let t be any target vertex. Placing $k - 1$ pebbles at every vertex except t results in a configuration that is not t -solvable, hence the pebbling number is larger than $(n - 1)(k - 1)$. If $(n - 1)(k - 1) + 1$ pebbles are distributed over the vertices such that t has 0 pebbles (otherwise the solution is trivial), then by the pigeonhole principle one vertex must have at least k pebbles, such that one pebble can be moved to t . \square

Trees

Now we turn to something a bit more impressive: a formula for the n -fold pebbling number of a vertex in a tree. This formula was proven by Chung, and uses a so called *maximum path-partition*.

Definition 2.11. Let T be a tree and let r be a vertex of T . Let T^* be a directed graph obtained from T by pointing all edges towards r . A path-partition rooted at r is a sequence of disjoint vertex sets V_1, V_2, \dots, V_n with a nonincreasing size, that each induce a directed path in T^* , such that $V_1 \cup V_2 \cup \dots \cup V_n = V(T) \setminus \{r\}$. A path-partition is maximum if the sequence of path sizes, $\#V_1, \#V_2, \dots, \#V_n$, majorizes that of any other path-partition (rooted at the same vertex) lexicographically.

Theorem 2.12. [3] Let $n \geq 1$ and $k \geq 2$. Let T be a tree and let r be a vertex of T . If V_1, V_2, \dots, V_m is a maximum path-partition of T rooted at r , then:

$$\pi_n(T^{(k)}, r) = n \cdot k^{\#V_1} + k^{\#V_2} + \dots + k^{\#V_m} - m + 1$$




Figure 3 – Maximum path-partition

Proof. We prove this by induction on the number of vertices of T . Suppose the theorem holds for all trees with fewer vertices than T . Remove r from T and obtain the subtrees T_1, \dots, T_s that T splits into, where T_1 contains the longest path directed towards r . For each $i \leq s$, let v_i be the vertex of T_i that is a neighbor of r . A maximum path-partition of T can be obtained from the union of, for every $i \leq s$, a maximum path-partition of T_i rooted at v_i with v_i added to its longest path.

Moving n pebbles to r can be achieved by any combination of moving n_i pebbles to v_i such that $\lfloor n_1/k \rfloor + \dots + \lfloor n_s/k \rfloor = n$. Note that in general, $\pi_{n+1}(G, v) - 1$ is the largest number of pebbles that can be distributed over G in such a way that it is impossible to move more than n pebbles to v . Using this insight we can see that every sequence n_1, \dots, n_s such that $\lfloor n_1/k \rfloor + \dots + \lfloor n_s/k \rfloor < n$ yields a lower bound for $\pi_n(T, r)$:

$$\pi_n(T, r) > \sum_{i \leq s} (\pi_{n_i+1}(T_i, v_i) - 1)$$

Let c be a configuration on T that is not n -fold r -solvable. Obtain a configuration c^* by removing all pebbles on r and adding them one of its neighbors. Note that c^* is also not n -fold r -solvable. Split c^* into c_1, c_2, \dots, c_s , where c_i is a configuration on T_i , and for every $i \leq s$, let n_i be the maximum number of pebbles that can be moved to v_i using the pebbles from c_i . Note that $|c| = |c^*| \leq \sum_{i \leq s} (\pi_{n_i+1}(T_i, v_i) - 1)$ and $\lfloor n_1/k \rfloor + \dots + \lfloor n_s/k \rfloor < n$. We conclude from this that the maximum of the above lower bounds is strict:

$$\pi_n(T, r) - 1 = \max \left\{ \sum_{i \leq s} (\pi_{n_i+1}(T_i, v_i) - 1) \mid \sum_{i \leq s} \lfloor n_i/k \rfloor < n \right\}$$

Observe that $k^a + k^b \geq k^{a-1} + k^{b+1}$ if $a > b$. By the induction hypothesis, each $\pi_{n_i+1}(T_i, v_i)$ can be replaced by an expression of the form $(n_i + 1)k^x + c$, where c is a constant and x is the size of the longest path in T_i (excluding v_i). Since T_1 contains the longest path, the maximum is achieved when $n_1 = nk - 1$ and $n_i = k - 1$ for $1 < i \leq s$. Using some extra work we can now rewrite the above expression into the desired form. \square

Diameter-2 Graphs

In a diameter-2 graph, any two vertices are either adjacent or have a common neighbor. The weight-2 pebbling number of such graphs was determined in [5] to be bounded by the number of vertices plus one. We present a simplified version of this proof. In the following, when it is written that *we may assume X*, it is implied that the case *not X* can easily be seen to be solvable. Note that *X* may actually yield a contradiction, in which case *not X* and the associated solution follow. In Figure 4, the gray area represents a part of the graph that is repeated, indexed by i .

Theorem 2.13. [5] For a simple diameter-2 graph G :

$$\pi(G^{(2)}) \leq \#V(G) + 1$$

Proof. Let c be a configuration on G such that $|c| = \#V(G) + 1$, and let t be a target vertex. We show that c is t -solvable. Label the vertices with more than one pebble x_1, x_2, \dots, x_m . We may assume no x_i is equal to t or a neighbor of t . For each x_i , pick a neighbor y_i that connects to t . We may assume that each y_i has zero pebbles, and that all y_i are distinct.




Figure 4

A simple calculation now implies that at least one vertex has 3 pebbles. We call this vertex v and remove it from the x_i vertices. We may assume that none of the x_i vertices are adjacent to v . For each x_i , pick a neighbor z_i that connects to v . We may assume that each z_i has zero pebbles, that all z_i are distinct, and that all z_i are distinct from all y_i . A simple calculation now implies that at least one vertex has four pebbles, such that the configuration is t -solvable. \square

Since the pebbling number is always at least the number of vertices, it follows from the previous result that, for every simple diameter-2 graph G , we have:

$$\pi(G^{(2)}) \in \{\#V(G), \#V(G) + 1\}$$

It is possible to characterize exactly which diameter-2 graphs have a pebbling number equal to their number of vertices. This was done in [6], where the terminology *Class-0* (for those simple graphs with a pebbling number equal to their number of vertices) and *Class-1* (for the other diameter-2 graphs) originated.

2.1 Questions

When Chung introduced pebbling numbers in her paper about hypercubes and the Erdős-Lemke conjecture, she also asked a number of general questions about them. It is worth revisiting those questions, since they are the primer to many later publications. Chung's definition of graph pebbling is analogous to ours when ignoring edge direction, *i.e.* making all edges bidirectional, and setting a standard edge weight of 2. Therefore, in the following questions, you should assume that all graphs are of this form.

Question 1. *Is it true that $\pi(G) = \max\{2^{\text{diam}(G)}, \#V(G)\}$?*

No. The star graph S_3 with three leaves has 4 vertices and a diameter of 2, but $\pi(S_3) = 5$. The 7-cycle C_7 has 7 vertices and a diameter of 3, but $\pi(C_7) = 11$.

Question 2. *Is it true that $\pi(G \square H) = \pi(G)\pi(H)$?*

No. Consider $C_3 \square P_3$. We have $\pi(C_3) = 3$ and $\pi(P_3) = 4$, while $\pi(C_3 \square P_3) = 9$. At the end of this chapter we will briefly explain a method based on CTL model checking that we implemented to automatically compute the pebbling number of $C_3 \square P_3$.

Question 3. (Ronald Graham) *Is it true that $\pi(G \square H) \leq \pi(G)\pi(H)$?*

Unsolved. This elusive question is often referred to as *Graham's conjecture*, after Ronald Graham who Chung attributes it to. A lot of graph pebbling research has been focused on this conjecture. The *2-pebbling property*, which was suggested by Chung based on her work about the pebbling number of the hypercube, has yielded quite some results.

Definition 2.14. *Let G be a graph. G has the 2-pebbling property (2PP) if for every configuration on G with at least $2\pi(G) - q + 1$ pebbles, where q equals the number of vertices that have at least one pebble, it is possible to move two pebbles to any target vertex.*

On many graphs, configurations that are more 'spread out' need, in the worst case, fewer pebbles to be solvable. Several authors have used the 2-pebbling property to prove special cases of Graham's conjecture. An overview of some major results is given below.

The Cartesian product of ...	Year	Reference
... K_n with a 2PP graph	1989	[3]
... a sequence of trees	1992	[4]
... two 5-cycles	1998	[7]
... a tree with a 2PP graph	2000	[8]
... a sequence of cycles (except C_5) and a 2PP graph	2008	[10]
... the Lemke graph ^a with a tree or with K_n	2017	[16]

Table 1 – Confirmed cases of Graham's conjecture.

^aSee Figure 5 on the next page.

Question 4. *Is it true that every graph has the 2-pebbling property?*

No. Paul Lemke identified the smallest graph that does not have the 2-pebbling property, which is now called the *Lemke graph* and is denoted as L . Two representations of the Lemke graph are shown below: the left depiction is standard in the literature, and the right depiction is the one we prefer to use. It can be determined that $\pi(L) = 8$. The configuration given in Figure 6 demonstrates that L does not satisfy the 2-pebbling property.




Figure 5 – Depictions of the Lemke graph




Figure 6 – L does not satisfy the 2PP.

The Lemke graph is interesting because it challenges the methods used to find the pebbling number of graph products. The pebbling number of $L \square L$ is still unknown; if there exists a configuration of 64 pebbles on this graph that is unsolvable for some target vertex, then Graham’s conjecture is false! Quite recently, with the help of linear optimization and so called *weight functions*, it was determined that $\pi(L \square L) \leq 85$ [19]. In Chapter 3 we explain the basic ideas underlying these weight functions.

Question 5. *Is $\pi(C_5 \square C_5 \square \dots \square C_5) = 5^n$, where we take the product of n 5-cycles?*

Unsolved. So far this equality has only been proven for C_5 and $C_5 \square C_5$ [7]. Among cycles, the 5-cycle is a curious exception, because Graham’s conjecture has been confirmed for products of cycles of *any* other size (see Table 1).

2.2 Computational Methods

Is it possible to determine pebbling numbers of moderately sized graphs using a computer? A brute-force approach is not going to work; even on relatively small graphs the number of possible configurations is very high. Let's ignore automorphisms for a moment, and calculate the total number of ways in which k pebbles can be distributed over n distinct vertices:

Theorem 2.15. *Let G be a graph with n vertices. For $k \in \mathbb{N}$:*

$$\#\{c : V(G) \rightarrow \mathbb{N} \mid |c| = k\} = \binom{k+n-1}{n-1}$$

Proof. Let v_1, v_2, \dots, v_n be the vertices of G . We assign k pebbles to these vertices. Pick numbers p_1, p_2, \dots, p_{n-1} from $\{1, 2, \dots, k+n-1\}$ such that $p_1 < p_2 < \dots < p_{n-1}$, this can be done in $\binom{k+n-1}{n-1}$ ways. Define a configuration c as follows:

$$c(v_i) := \begin{cases} p_1 - 1 & \text{if } i = 1 \\ p_i - (p_{i-1} + 1) & \text{if } 1 < i < n \\ (k+n) - (p_{i-1} + 1) & \text{if } i = n \end{cases}$$

Now $|c| = k$. Note that every possible configuration of size k can be picked in this way. \square

Two configurations that only differ by an automorphism permuting the vertices admit the same solutions. We could divide the above formula by the number of automorphisms to get a rough estimate of the number of *different* configurations. By this metric there are *at least* 770 different configurations of 10 pebbles on the Petersen graph, and more than 3×10^{11} different configurations of 25 pebbles on $C_5 \square C_5$ ⁱⁱ. To date, no general algorithm for pebbling numbers has been found that can handle $C_5 \square C_5$. It is worth noting that just solving individual configurations is already NP-complete [9]. At the end of Chapter 4 we describe a general method to find pebbling solutions using integer programming.

Barely Sufficient Configurations In [11] a method is developed to find the set of *barely sufficient* configurations using a recursive algorithm that applies inverse pebbling steps. A configuration is barely sufficient when it is solvable for some vertex, and it is no longer solvable when one pebble is removed from any vertex. The weight-2 pebbling number of all simple graphs with fewer than 10 vertices was computed using this method (combined with a graph simplification technique) in about a day. This publication is also the only one we found that describes graph pebbling using weighted graphs.

CTL model checking We implemented a method to compute pebbling numbers of small graphs using *CTL model checking* and the NuSMV software. A CTL model describes, over a finite space of variables; a set of initial states, a state transition relation, and a specification about reachable states. The specification is verified by the model checker using a highly optimized algorithm based on binary decision diagrams. Both the set of initial states and the transition relation can be expressed as a boolean formula over the model variables.

ⁱⁱThe graph $C_5 \times C_5$ has 200 automorphisms.

```

MODULE main
DEFINE n := 4;
VAR c : array 1..3 of 0..n;
INIT c[1] + c[2] + c[3] = n

TRANS
( c[1]>1 & next(c[1])=c[1]-2 & next(c[2])=c[2]+1 & next(c[3])=c[3] ) |
( c[2]>1 & next(c[1])=c[1]+1 & next(c[2])=c[2]-2 & next(c[3])=c[3] ) |
( c[2]>1 & next(c[1])=c[1] & next(c[2])=c[2]-2 & next(c[3])=c[3]+1 ) |
( c[3]>1 & next(c[1])=c[1] & next(c[2])=c[2]+1 & next(c[3])=c[3]-2 ) |
( next(c[1])=c[1] & next(c[2])=c[2] & next(c[3])=c[3] )

SPEC EF c[1] > 0
SPEC EF c[2] > 0
SPEC EF c[3] > 0

```

Figure 7 – NuSMV model for P_3

The above figure shows a NuSMV model for weight-2 pebbling on P_3 . We want to check that starting at any configuration of four pebbles on P_3 , it is possible to move one pebble to any target vertex. The variables $c[1]$, $c[2]$ and $c[3]$ express the number of pebbles on each vertex. Since there can at most be four pebbles on any vertex, the value of these variables ranges from zero to four. The INIT formula describing the initial set of states is true for any configuration of four pebbles. The TRANS formula describing the transition relation between the current state, described by $c[1]$, $c[2]$ and $c[3]$, and the next state, described by $\text{next}(c[1])$, $\text{next}(c[2])$ and $\text{next}(c[3])$, is true for all valid pebbling steps and, for technical reasons, a transition where no pebbles are moved. For each vertex, a formula is defined using SPEC to specify that one pebble can be moved to that vertex from any initial state. The specification formula EF φ is true if, for all initial states, there is a sequence of transitions such that the formula φ is true in the final stateⁱⁱⁱ.

Figure 8 on the next page shows a part of a NuSMV model for the Lemke graph. Here the initial states contain all configurations from which, to satisfy the 2-pebbling property, it should be possible to move two pebbles to any target vertex. When verifying this model, NuSMV finds a counterexample to SPEC EF $c[1] > 0$, such as the one shown in Figure 6. It takes NuSMV a few seconds to verify the weight-2 pebbling number of graphs like $P_3 \square C_3$ and L , and doing so for the Petersen graph already takes around 10 seconds. Larger graphs quickly become intractable.

ⁱⁱⁱSee page 37 of [13].

2 Pebbling Numbers

```
MODULE main
DEFINE n := 8; p := 8;
VAR c : array 1..n of 0..2*p;

INIT
  c[1] + c[2] + c[3] + c[4] + c[5] + c[6] + c[7] + c[8] = 2*p + 1 -
  count(c[1]>0, c[2]>0, c[3]>0, c[4]>0, c[5]>0, c[6]>0, c[7]>0, c[8]>0)

TRANS
  ( c[1]>1 & next(c[1])=c[1]-2 & next(c[2])=c[2]+1 & ... ) |
  ( c[1]>1 & ... ) |
  ...

SPEC EF c[1] > 1
SPEC EF c[2] > 1
SPEC EF c[3] > 1
...
```

Figure 8 – NuSMV model to check if L satisfies the 2PP.

```
-- specification EF c[1] > 1 is false
-- as demonstrated by the following execution sequence
Trace Description: CTL Counterexample
Trace Type: Counterexample
-> State: 1.1 <-
  c[1] = 0
  c[2] = 0
  c[3] = 0
  c[4] = 1
  c[5] = 1
  c[6] = 1
  c[7] = 1
  c[8] = 8
  p = 8
  n = 8
-- specification EF c[2] > 1 is true
-- specification EF c[3] > 1 is true
...
```

Figure 9 – NuSMV output

3 Weight Functions

In this chapter we use weight functions that assign weights to the *vertices* of a graph. These weight functions should be distinguished from the edge-weight functions we introduced earlier. Weight functions enable us to recognize solvable pebbling configurations with linear inequalities, offering a powerful method for finding pebbling numbers. The use of weight functions and linear optimization was proposed by Hurlbert [12][15]. We only discuss weight functions applied to weight-2 pebbling, so for the duration of this chapter you should assume that the edge weight is always 2. We start by proving a key theorem, and showing how the pebbling number of various graphs can be determined using a set of weight functions.

Definition 3.1. *Let G be a graph and t one of its vertices. A function $w : V(G) \rightarrow \mathbb{Q}^+$ is a weight function for t if $w(t) = 0$, and for every vertex $u \in \text{supp}(w)$ such that $(u, t) \notin E(G)$, there is an edge (u, v) such that $w(v) \geq 2 \cdot w(u)$.ⁱ*

Given a weight function, we can compute the weight of a configuration by multiplying, for each vertex, the number of pebbles with the weight assigned by the weight function, and adding the results. For a weight function w and a configuration c this is denoted as $|w \cdot c|$. The defining property of a weight function makes sure that pebbles on a positively weighted vertex not adjacent to the target vertex can be moved to a neighboring vertex (via pebbling steps) without the configuration as a whole losing weight. This way we can prove the following result, which Hurlbert calls the *Weight Function Lemma*.

Lemma 3.2. *Let G be a graph and t one of its vertices. If w is a weight function for t , and c is a configuration on G such that $|w \cdot c| > |w|$, then c is t -solvable.*

Proof. Let c be a configuration on G such that $|w \cdot c| > |w|$. Determine a vertex u with $c(u) \geq 2$. If $(u, t) \in E(G)$, then we can move one pebble to t , and so c is t -solvable. Otherwise, determine an edge (u, v) such that $w(v) \geq 2 \cdot w(u)$, and apply one pebbling step to this edge, removing two pebbles from u and adding one to v . For the updated configuration c' we have $|w \cdot c'| \geq |w \cdot c| > |w|$. Repeat this procedure until a pebble has been moved to t . This should only take a finite number of steps, since there is only a finite number of pebbles. We conclude that c is t -solvable. \square

We are going to use sums of weight functions to determine the pebbling number of several graphs, including the Petersen graph and cycle graphs. Suppose a sequence w_1, w_2, \dots, w_n of weight functions for a vertex t in a graph G is given. If the sum $w = w_1 + w_2 + \dots + w_n$ is positive for all vertices except t , then we say that w covers G , and if all positive values of w are equal, then we say that w covers G uniformly. The smallest positive value of w is called its *minimum weight*.

ⁱIn accordance with the literature, weight values are non-negative rationals. This can be useful since the neighbors of the target vertex could just as well be given a weight of 1, their neighbors a weight of $\frac{1}{2}$, and so on.

3 Weight Functions

Lemma 3.3. Let G be a graph and t one of its vertices. Let w_1, w_2, \dots, w_n be weight functions for t , and let $w = w_1 + w_2 + \dots + w_n$. If w covers G and m is its minimum weight, then:


$$\pi(G, t) \leq \lfloor |w|/m \rfloor + 1$$

Proof. Let c be a configuration on G with $|c| \geq \lfloor |w|/m \rfloor + 1$. We prove that c is t -solvable. Suppose t has no pebbles, i.e. $c(t) = 0$. Now $|w \cdot c| \geq |c| \cdot m > |w|$, which we can expand into $|w_1 \cdot c| + \dots + |w_n \cdot c| > |w_1| + \dots + |w_n|$. Determine $i \leq n$ such that $|w_i \cdot c| > |w_i|$ and apply Lemma 3.2 to conclude that c is indeed t -solvable. \square

Theorem 3.4. For the Petersen graph P holds:

$$\pi(P) = 10$$

Proof. Since P has 10 vertices, $\pi(P) \geq 10$. Note that all vertices of P have the same pebbling number. Pick some target vertex t . For each neighbor of t there is a weight function like the one shown below. Adding these three weight functions together covers P uniformly with fours. Using Lemma 3.3 we find $\pi(P, t) \leq \lfloor 36/4 \rfloor + 1 = 10$.




\square

Theorem 3.5. For the complete bipartite graph $K_{m,n}$ with $m, n \geq 2$ holds:

$$\pi(K_{m,n}) = m + n$$

Proof. The sum of the two weight functions shown below covers the graph uniformly with twos. Using Lemma 3.3 we find $\pi(K_{m,n}) \leq m + n$.




\square

Theorem 3.6. For the Lemke graph L with v_3 as in Figure 5 holds:

$$\pi(L, v_3) = 8$$

Proof. Use the three weight functions shown below.




□

Theorem 3.7. For the even cycle C_{2n} with $n \geq 1$ holds:

$$\pi(C_{2n}) = 2^n$$

Proof. Make two mirror copies of the weight function shown below, and call their sum w . The minimal weight of w is 2 and $|w| = 2(2^n - 1)$. Using Lemma 3.3 we find $\pi(C_{2n}) \leq 2^n$. This is strict, since the configuration that puts $2^n - 1$ pebbles on the vertex furthest from the target is unsolvable.




□

Theorem 3.8. For the odd cycle C_{2n+1} with $n \geq 1$ holds:

$$\pi(C_{2n+1}) = 2\lfloor 2^{n+1}/3 \rfloor + 1$$

Proof. Make two mirror copies of the weight function shown below, and call their sum w . The minimal weight of w is 3 and $|w| = 2(2^{n+1} - 1)$. It follows from Lemma 3.3 that $\pi(C_{2n+1}) \leq \lfloor 2(2^{n+1}/3 - 1/3) \rfloor + 1 = 2\lfloor 2^{n+1}/3 \rfloor + 1$, where the last equality can be determined by analyzing the possible remainders of $2^{n+1}/3$. To see that this bound is strict, put $\lfloor 2^{n+1}/3 \rfloor$ pebbles on each of the two vertices furthest from the target, and note that no solution is possible via either side of the cycle since $\lfloor 2^{n+1}/3 \rfloor + \lfloor 2^n/3 \rfloor < 2^n$.



□

3.1 Limitations

So far weight functions have been quite instrumental, but there are some serious limitations. For some configurations there is no weight function to solve it using the procedure in the proof of Lemma 3.2. If for a configuration c there is a weight function w with $|w \cdot c| > |w|$, then we say that c can be solved using a weight function, and if every configuration c on a graph G with $|c| = \pi(G)$ can be solved using a weight function, then we say that G can be solved using weight functions. In this section we look at several graphs that are *not* solvable using weight functions.

Splitting Structures

Consider the configurations on the cube and the Lemke graph in the figure below. Both are solvable, but in order to solve them the heap of 5 pebbles must be divided over two neighbors to take advantage of the extra pebbles which are placed there. There is no other solution. Hurlbert calls these situations *splitting structures*.




Figure 10

In the procedure that we used to prove Lemma 3.2 it does not matter how pebbles are distributed over neighbors, as long as the configuration weight is not decreased. In both of the above configurations, it is impossible to come up with a weight function to solve it. Both configurations also have the same number of pebbles as the (weight-2) pebbling number of their respective graphs. If the pebbling number of either of these graphs could be determined using just weight functions, then there should also be a weight function to solve these specific configurations. Hence we are forced conclude that the pebbling number of neither the cube nor the Lemke graph can be found using weight functions alone.

In [14] this issue is bypassed by introducing a special weight function called a ‘lollipop’ that contains a path attached to an even cycle. On a lollipop, the vertex furthest from the target gets slightly more weight than it can have in an ordinary weight function. This breaks the simple proof of Lemma 3.2, since the configuration weight *decreases* when a pebbling step removes pebbles from this vertex.

But in Lemma 3 of [14] the authors show that their lollipop weight function w_l has the same property as ordinary weight functions: If $|w_l \cdot c| > |w_l|$, then c is solvable. An example of a lollipop weight function on the cube, solving the problematic configuration from Figure 10, is shown below.




Figure 11 – A lollipop weight function

Balanced Configurations

Consider the two configurations on S_3 below; the left one is solvable and the right one is not. Since the left configuration is the average of two unsolvable configurations like the right one, it cannot be solved using a weight function.




Figure 12

This example has fewer than $\pi(S_3)$ pebbles, and it is still possible to show that $\pi(S_3) = 5$ using weight functions. This led us to ask the following question: *Given a graph G , can we find $\pi(G)$ using just weight functions if all configurations c with $|c| = \pi(G)$ can be solved without splitting?*

It appears the answer is *no*. It has been shown that $\pi(C_5 \square C_5) = 25$, and as far as we know all configurations of 25 pebbles on $C_5 \square C_5$ can be solved without splittingⁱⁱ. But we did find a configuration of 25 pebbles that is not solvable using a weight function. The offending configuration, and a weight function that comes closest to solving it, are shown in Figure 13.

ⁱⁱA careful analysis of [7] might be able to confirm this.

3 Weight Functions




Figure 13 – A configuration (left) and a weight function (right) on $C_5 \square C_5$.

3.2 Linear Programming

Every weight function w gives rise to a linear inequality of the form $|w \cdot c| \leq |w|$ that holds if a configuration c cannot be solved with it. Linear programming software can efficiently determine the optimal solution of a linear objective function subject to linear constraints. When we use the configuration size as the objective function, and the linear inequalities belonging to a set of weight functions as constraints, then the maximal solution gives the size of the largest possible configuration for which none of the weight functions work; all larger configurations will violate one of the weight function constraints, and are therefore solvable.

Fractional Relaxation A pebbling configuration consists of integer amounts; we do not split up pebbles between vertices, and thus we need an LP solver that finds integer solutions. In linear programming, finding integer solutions is much more computationally intensive than finding fractional solutions, and therefore our problem becomes much easier to compute if we ‘relax’ it to configurations with fractional pebble amounts. Since integer configurations are a subset of fractional configurations, this will still yield a useful bound.

Lemma 3.9. *Let G be a graph and t one of its vertices. If w_1, w_2, \dots, w_n are weight functions for t , and a fractional configuration $c_{\max} : V(G) \rightarrow \mathbb{Q}^+$ maximizes $|c_{\max}|$ while satisfying $|w_i \cdot c_{\max}| \leq |w_i|$ for all $i \leq n$, then:*

$$\pi(G, t) \leq \lfloor |c_{\max}| \rfloor + 1$$

Proof. Let c be a configuration on G with $|c| \geq \lfloor |c_{\max}| \rfloor + 1$. Then $|c| > |c_{\max}|$. Because c_{\max} is maximal, there must exist an $i \leq n$ such that $|w_i \cdot c| > |w_i|$. It now follows from Lemma 3.2 that c is t -solvable. \square

Dual Solutions The *dual solution* of an LP solver reveals the constraints by which the optimal corner of the feasibility polytopeⁱⁱⁱ is bounded. It is possible to enter many weight functions into an LP solver, and then use the dual solution to find out which weight functions are sufficient to obtain the optimal bound. Hurlbert did this in [12]; he first picked a random graph on 15 vertices, then a program generated over 20 thousands weight functions for the optimization problem, and finally the CPLEX solver found an optimum based on just 11 of these weight functions. In this method, the feasibility polytope for a graph with n vertices is n -dimensional, where each dimension represents the number of pebbles on one of its vertices. The corners of a 3-dimensional polytope are almost always at the intersection of 3 of its faces, and in general a linear optimization problem with n variables typically requires at most n constraints in its certificate.




Figure 14 – Integer solutions

ⁱⁱⁱThe feasibility polytope is the region containing all points that satisfy the given linear constraints.

4 Pebble Flows

What can we say about the direction in which pebbles are moved? Hurlbert calls solutions *greedy* when pebbles are moved strictly closer to the target vertex. For some graphs, for example for cubes, there is a greedy solution for all configurations with a size equal to the graph's pebbling number. But sometimes pebbles actually have to be moved *away* from the target vertex. This is the case in the weight-2 example shown below. It is possible to determine that the weight-2 pebbling number of this graph is 10 using weight functions.




Figure 15

Do we ever need to move pebbles back and forth between two vertices? This does not seem very useful; if pebbles that are moved to some vertex can be returned later, then how are they essential to a solution? In this chapter we show that, indeed, to solve any particular configuration it is never needed to move pebbles back-and-forth between any two vertices. To prove this, we regard solutions as *flow networks*. We believe that this approach might also be useful for other graph pebbling problems.

Definition 4.1. Let G be a graph. A pebble flow \mathcal{F} on G is a pair (c, f) of a configuration c on G , and a flow function $f : V(G) \times V(G) \rightarrow \mathbb{N}$. The value $f(u, v)$ represents the number of pebbles that are added to v using pebbling steps via the edge (u, v) . If $(u, v) \notin E(G)$ then $f(u, v) = 0$. The configuration of \mathcal{F} is also written as $c_{\mathcal{F}}$ and the flow from u to v as $\mathcal{F}(u, v)$.

A pebble flow combines a configuration with a solution expressed in flow values. We can compute several useful values from this, including the total inflow and outflow of a vertex, the *weighted* outflow of a vertex which is the number of pebbles that actually need to be removed to realize its outflow, and the *excess* of a vertex which is the number of pebbles that are left after the inflow and outflow are satisfied. We will define these terms more precisely on the next page.

4 Pebble Flows

Definition 4.2. Let G be a graph. For a pebble flow \mathcal{F} on G and a vertex $v \in V(G)$, the inflow $\text{in}_{\mathcal{F}}(v)$, outflow $\text{out}_{\mathcal{F}}(v)$, weighted outflow $\text{out}_{\mathcal{F}}^*(v)$, and excess $x_{\mathcal{F}}(v)$ of v are defined as:

$$\begin{aligned}\text{in}_{\mathcal{F}}(v) &:= \sum_{u \in V(G)} \mathcal{F}(u, v) \\ \text{out}_{\mathcal{F}}(v) &:= \sum_{u \in V(G)} \mathcal{F}(v, u) \\ \text{out}_{\mathcal{F}}^*(v) &:= \sum_{u \in V(G)} \omega_G(v, u) \cdot \mathcal{F}(v, u) \\ x_{\mathcal{F}}(v) &:= c_{\mathcal{F}}(v) + \text{in}_{\mathcal{F}}(v) - \text{out}_{\mathcal{F}}^*(v)\end{aligned}$$

We extend the step relation for pebbling configurations to pebble flows as follows:

Definition 4.3. Let G be a graph. The pebble flow step relation $\mathcal{F}_1 \rightarrow \mathcal{F}_2$ between two pebble flows \mathcal{F}_1 and \mathcal{F}_2 on G holds if there is an edge (u, v) such that $\mathcal{F}_1(u, v) \geq 1$ and $c_{\mathcal{F}_1}(u) \geq k$ for $k = \omega_G(u, v)$, and:

$$\begin{aligned}\forall w : c_{\mathcal{F}_2}(w) &= \begin{cases} c_{\mathcal{F}_1}(u) - k & \text{if } w = u \\ c_{\mathcal{F}_1}(v) + 1 & \text{if } w = v \\ c_{\mathcal{F}_1}(w) & \text{otherwise} \end{cases} \\ \forall x, y : \mathcal{F}_2(x, y) &= \begin{cases} \mathcal{F}_1(u, v) - 1 & \text{if } (x, y) = (u, v) \\ \mathcal{F}_1(x, y) & \text{otherwise} \end{cases}\end{aligned}$$

4.1 Properties

To translate a step-by-step solution into a pebble flow, you have to count the number of pebbling steps along every edge. For every vertex, there will never flow more pebbles out (the weighted outflow) than the number of pebbles put there by the initial configuration plus the number of pebbles flowing in. This means that the excess of every vertex is non-negative. We call this property *feasible*. Clearly, a solvable configuration corresponds to a feasible pebble flow.

Definition 4.4. A pebble flow \mathcal{F} is feasible if $\forall v : x_{\mathcal{F}}(v) \geq 0$.

Lemma 4.5. Let G be a graph and t one of its vertices. If c is an n -fold t -solvable configuration, then there exists a feasible pebble flow \mathcal{F} with $c_{\mathcal{F}} = c$ and $x_{\mathcal{F}}(t) \geq n$.

Proof. Left to the reader. □

The above property is rather trivial, but can we turn it around? Does every feasible pebble flow correspond to a step-by-step solution? If we can prove that this is true, then we are very close to answering the question about back-and-forth movement between vertices.

Realization

We quickly run into the following problem: it is possible to create a kind of deadlock using a cycle. This is illustrated below, where each edge has a weight of 2 and a flow value of 1. While this pebble flow is feasible, it is *not* possible to apply further pebbling steps and reduce all flow values to zero.




Figure 16 – Cyclic flow

However, in a situation like this there isn't really anything to gain. The excess of all vertices is zero, and each vertex already has more than zero pebbles, so why not stop here? That is exactly what we are going to do: we will show that for any pebble flow it is possible to take steps forward until every vertex has at least as many pebbles as its excess, at which point we say that the pebble flow is *realized*.

Definition 4.6. A pebble flow \mathcal{F} is realized if $\forall v : c_{\mathcal{F}}(v) \geq x_{\mathcal{F}}(v)$.

Definition 4.7. A pebble flow \mathcal{F} is realizable if $\mathcal{F} \rightarrow^* \mathcal{F}^*$ for some \mathcal{F}^* that is realized.

We found two methods to prove that all feasible pebble flows are realizable. The first method is to remove all flow cycles, leaving an acyclic pebble flow that can be realized starting at its leaves. Note that such a pebble flow will always have at least one leaf; a vertex with no inflow, and that, since it is feasible, this vertex has enough pebbles to fully realize its outflow. The second method we found doesn't need to know anything about flow cycles; it just shows that, as long as the pebble flow is not fully realized, there must be a vertex somewhere with enough pebbles for a next pebbling step. Since there can only be a finite number of steps until all pebbles run out, a realized configuration should eventually be reached. We will work out the second method in more detail.

Lemma 4.8. For a pebble flow \mathcal{F} on G , the following holds:

$$\sum_{v \in V(G)} \text{in}_{\mathcal{F}}(v) = \sum_{v \in V(G)} \text{out}_{\mathcal{F}}(v)$$

Proof. Every flow value counts as inflow for one vertex and outflow for another vertex. Thus the sum of the inflow of all vertices is equal to the sum of the outflow of all vertices. \square

Lemma 4.9. *Every feasible pebble flow is realizable.*

Proof. Suppose \mathcal{F} is a pebble flow on a graph G , and \mathcal{F} is not realized. Determine a vertex v such that $c_{\mathcal{F}}(v) < x_{\mathcal{F}}(v)$. Since the excess has not been realized, v must still be waiting for more pebbles than it needs to send away: $\text{in}_{\mathcal{F}}(v) > \text{out}_{\mathcal{F}}^*(v) \geq \text{out}_{\mathcal{F}}(v)$. It follows from Lemma 4.8 that there also exists a vertex in G where this inequality is reversed. Determine w such that $\text{in}_{\mathcal{F}}(w) < \text{out}_{\mathcal{F}}(w)$. Let k be the minimum weight of an outflow of w , such that:

$$\text{out}_{\mathcal{F}}^*(w) \geq k \cdot \text{out}_{\mathcal{F}}(w) \geq k \cdot (\text{in}_{\mathcal{F}}(w) + 1) \geq \text{in}_{\mathcal{F}}(w) + k$$

It follows from the feasibility of \mathcal{F} that $c(w) \geq k$, and thus that w has enough pebbles for one pebbling step. Repeat this process until the pebble flow is realized. Note that, as long as it is not realized, there must be enough pebbles somewhere to apply a pebbling step, and that we can only repeat this a finite number of times because every step reduces the total number of pebbles. \square

Unidirectional

We have shown that pebble flows correspond to step-by-step solutions, and step-by-step solutions correspond to pebble flows. To answer our question about back-and-forth movement it only remains to see that any solution gives rise to a solution without back-and-forth movement. We call such a solution *unidirectional*.

Theorem 4.10. *Let G be a graph. If a configuration c is n -fold t -solvable, then there is a unidirectional solution.*

Proof. Suppose c is n -fold t -solvable. Determine a feasible pebble flow \mathcal{F} corresponding to a solution of c , such that $x_{\mathcal{F}}(t) \geq n$. We define a new pebble flow \mathcal{F}^* with the same configuration that only describes the net flow between any two vertices:

$$\mathcal{F}^*(u, v) := \mathcal{F}(u, v) - \min\{\mathcal{F}(u, v), \mathcal{F}(v, u)\}$$

This pebble flow is feasible *and* unidirectional. The excess does not decrease at any vertex, and in particular we still have $x_{\mathcal{F}^*}(t) \geq n$. Using Lemma 4.9 we can find a step-by-step solution corresponding to \mathcal{F}^* that puts at least n pebbles on t . In this solution there is no back-and-forth movement. \square

4.2 Other Applications

Transforming a pebbling solution into a flow network removes the temporal aspect of steps, and allows the solution to be presented *at once*. The feasibility property forms a conjunction of linear inequalities; an integer programming problem! This provides a straightforward method to solve pebbling configurations.

Optimal Pebbling The *optimal* pebbling number of a graph is the minimum size of a configuration that is solvable for all of its vertices, *i.e.* such that one pebble can be moved to any desired vertex. There are several unresolved questions about the optimal pebbling number. Curiously, the weight-2 optimal pebbling number of square grids is not yet known. A lower bound was established quite recently in [18]:

$$\pi_{\text{opt}}(P_m^{(2)} \square P_n^{(2)}) \geq \frac{2}{13}mn$$

An asymptotic upper bound is obtained in [17] using a configuration on an infinite grid that, on average, spreads every pebble over 3.5 vertices (the *covering ratio*). This configuration is shown in the figure below.




Figure 17 – Configuration with covering ratio 3.5.

Optimal configurations via integer programming By expressing solutions as pebble flows, finding an optimal configuration can be achieved with integer programming. We tried using the Z3 solverⁱ to find optimal configurations for small grids, but this already appeared to be intractableⁱⁱ for grids larger than 3×4 . Perhaps it is possible to use smarter solving heuristics to find optimal configurations for larger grids using this method. For example, the distance over which pebbles are moved across the grid could be restricted. This problem might also be an interesting benchmark for LP solvers.

ⁱ<https://www.microsoft.com/en-us/research/project/z3-3/>

ⁱⁱComputing a solution for a 3×4 grid using the Z3 solver took about 10 seconds on an AMD Ryzen 7 5700G CPU. We ran a similar calculation for a 4×4 grid, but we terminated it after about ten minutes.

5 Hypercubes

We now turn to the problem that initiated the study of graph pebbling: the pebbling number of the hypercube. Chung's method to prove the Erdős-Lemke conjecture, which we stated in the introduction, uses the pebbling number of a hypercube with arbitrary edge weights under the restriction that parallel edges have the same weight. Given the weights k_1, k_2, \dots, k_n , this corresponds to the following n -dimensional hypercube:

$$P_2^{(k_1)} \square P_2^{(k_2)} \square \dots \square P_2^{(k_n)}$$

Chung determined that the pebbling number of this graph is the product of the weights; $k_1 \cdot k_2 \cdot \dots \cdot k_n$, but the proof she presents leaves some details to the reader, and we actually got stuck trying to work it out. We filled this gap by developing a slightly different proof, which is presented in this chapter.

5.1 Analysis

We start by analyzing Chung's proof in [3] that an n -dimensional hypercube with a standard edge weight of 2 has pebbling number 2^n . We have included a copy of this proof below in its original form, with some minor adjustments.

Theorem 5.1. [3] *In an n -cube with a specified vertex v , the following are true:*

- (i) *If 2^n pebbles are assigned to vertices of the n -cube, one pebble can be moved to v .*
- (ii) *Let q be the number of vertices that are assigned an odd number of pebbles. If there are in total more than $2^{n+1} - q$ pebbles, then two pebbles can be moved to v .*

Proof. The proof is by induction on n . It is trivially true for $n = 0$. Suppose it is true for n , we will prove it for $n + 1$. Partition the $(n + 1)$ -cube into two n -cubes called M_1 and M_2 , where v is in M_1 . Let v' denote the vertex in M_2 adjacent to v . The edges between M_1 and M_2 form a perfect matching. Suppose M_i contains p_i pebbles with q_i vertices having an odd number of pebbles, for $i = 1, 2$.

(i) Suppose there are $p \geq 2^{n+1}$ pebbles assigned to vertices of the $(n + 1)$ -cube. If $p_1 \geq 2^n$, then by induction one pebble can be moved to v in M_1 . We may assume $p_1 < 2^n$ and we consider the following two cases:

Case (a1) $q_2 > p_1$. Since $p_2 = p - p_1 > 2^{n+1} - q_2$, by induction from (ii) in M_2 two pebbles can be moved to v' . Therefore one pebble can be moved to v .

Case (a2) $q_2 \leq p_1$. We can move at least $(p_2 - q_2)/2$ pebbles from M_2 to M_1 . This results in a total of $p_1 + (p_2 - q_2)/2 \geq p_1 + (p_2 - p_1)/2 = (p_1 + p_2)/2 \geq 2^n$ pebbles in M_1 . By induction, we can then move one pebble to v .

5 Hypercubes

(ii) Suppose there are $p = p_1 + p_2 > 2^{n+2} - q_1 - q_2$ pebbles assigned to vertices of the $(n+1)$ -cube. We want to show that two pebbles can be moved to v . We consider three cases:

Case (b1) $p_1 > 2^{n+1} - q_1$. By induction from (ii) in M_1 , two pebbles can be moved to v .

Case (b2) $2^n \leq p_1 \leq 2^{n+1} - q_1$. Since $p_1 \geq 2^n$, by induction from (i) in M_1 one pebble can be moved to v . Note that $p_2 = p - p_1 > 2^{n+2} - q_1 - q_2 - p_1 \geq 2^{n+1} - q_2$. Therefore, by induction from (ii) in M_2 , two pebbles can be moved to v' , such that a second pebble can be moved to v .

Case (b3) $p_1 < 2^n$. For any integer t satisfying $p_2 \geq q_2 + 2t$, it is possible to move t pebbles to M_1 while $p_2 - 2t$ pebbles remain in M_2 . Note:

$$p_2 > 2^{n+2} - q_1 - q_2 - p_1 = (2^{n+1} - q_2) + (2^{n+1} - q_1 - p_1) \geq q_2 + (2^{n+1} - q_1 - p_1)$$

The last inequality follows since q_2 is at most the number of vertices in M_2 . Now taking t to be $2^n - (p_1 + q_1)/2$, we move t pebbles to M_1 , leaving more than $2^{n+1} - q_2$ pebbles in M_2 . In M_1 there are now $p_1 + 2^n - (p_1 + q_1)/2 = 2^n + (p_1 - q_1)/2 \geq 2^n$ pebbles. We can then move one pebble to v in M_1 and move two pebbles to v' in M_2 , enabling us to move a second pebble to v . \square

Odd 2-Pebbling Property

Where does property (ii) in Theorem 5.1 come from? One way to look at it is like this; if in M_1 we want to move one pebble to v , then the first thing we could try is moving as many pebbles from M_2 to M_1 as we can, and seeing if this results in enough pebbles on M_1 to apply the induction hypothesis. The number of pebbles remaining in M_2 will be equal to q_2 ; the number of vertices that have an odd number of pebbles. This is exactly what is happening in case (a2) of the proof. The above strategy works as long as q_2 does not exceed the total number of pebbles on M_1 . If it does, then M_2 has enough pebbles to apply (ii), as is done in case (a1). Property (ii) is a kind of ‘residual case’ resulting from the above strategy, and it is sometimes called the *odd 2-pebbling property*.

When generalizing the weight of the edges between M_2 and M_1 , the number of pebbles that are left behind on M_2 in case (a2) is no longer equal to q_2 . If the weight between M_2 and M_1 is k , then in the worst case $k - 1$ pebbles are left behind on *every* vertex. We are going to ‘upgrade’ the proof of Theorem 5.1 to hypercubes with varying weights using an alternative for property (ii).

Pebble Extraction

Could we prove Theorem 5.1 when (ii) is replaced with the normal 2-pebbling property? Going forward we will use $s^\#(c)$ to denote $\#\text{supp}(c)$, *i.e.* the number of vertices on which the configuration c puts at least one pebble. Let c be a configuration on the n -cube, can we prove the following, using the same case distinctions and strategies as the proof of Theorem 5.1?

$$\text{If } |c| \geq 2^{n+1} - s^\#(c) + 1, \text{ then } c \text{ is 2-fold } v\text{-solvable.}$$

The answer is *no*, as is demonstrated by the following example on the cube denoted by Q_3 . Let Q_3 be composed of the squares M_1 and M_2 with corresponding configurations c_1 and c_2 as shown below. Note that $\pi(C_4) = 4$ and $\pi(Q_3) = 8$. The following holds:

$$|c_1| + |c_2| = 10 \geq 2\pi(Q_3) - s^\#(c_1) - s^\#(c_2) + 1$$




Figure 18

According to the 2-pebbling property we should now be able to move two pebbles to any vertex in M_1 . We are trying to follow the same steps as Chung’s proof, so we apply case (b3) since $|c_1| = 3 < \pi(C_4)$. According to (b3) we should first move one pebble to M_1 using two pebbles from M_2 , and then apply the 2-pebbling property inductively in M_2 . But removing two pebbles from a vertex of M_2 *decreases* the support of c_2 , and the resulting configuration no longer satisfies the condition needed to apply the 2-pebbling property. We have to do more to generalize the proof.

Chung adds *pebble extraction* to the 2-pebbling property to solve this: If a configuration c on the n -cube contains at least $2^{n+1} - s^\#(c) + 1 + m$ pebbles, then it is possible to move two pebbles to a target vertex *and also* move $\lfloor m/k \rfloor$ pebblesⁱ to an adjacent n -cube using pebbling steps of weight k . Note that the latter is *not* possible with every configuration of just m pebbles, so there must be some interaction between these two parts.

We encountered a restriction on the weight of these extra pebbling steps that Chung does not mention. She proves that it works for $k = 2$ if all edges have a weight of 2, but when generalizing k we run into a problem. Let c_3 be the configuration below. For $m = 4$ we have:

$$|c_3| \geq 2\pi(C_4) - s^\#(c_3) + 1 + m$$




Figure 19

If we pick $k = 4$ then $\lfloor m/k \rfloor = 1$, and according to the above extraction property we should be able to find a vertex from which we can remove 4 pebbles. But c_3 does not have such a vertex, so it is not possible to move one pebble to an adjacent n -cube connected by edges of weight 4. To solve this problem we will apply induction from the dimension with the largest to the dimension with the smallest weight.

ⁱTheorem 3 in [3] actually allows a combination of pebbling steps with different weights.

Reduced Size

Our alternative to Chung's pebble extraction property was inspired by a detail in her proof. According to her proof, it is possible to extract at least $\lceil (|c| - s^\#(c))/2 \rceil$ pebbles from any configuration c using pebbling steps of weight 2. How do we prove this, and how does it generalize to other weights? Extracting n pebbles at a weight of k is similar to finding n blocks (pebbles placed on the same vertex) of k pebbles. In the following theorem, we prove a lower bound for finding these blocks.

Definition 5.2. Let V be a finite set. Let $c : V \rightarrow \mathbb{N}$ and $n, k \in \mathbb{N}$. We say that n blocks of k can be extracted from c if there is a function $e : V \rightarrow \mathbb{N}$ indicating the number of blocks that can be extracted from each vertex, such that $|e| = n$ and $c(v) \geq k \cdot e(v)$ for all $v \in V$. The result of this extraction is a configuration c' defined as $c'(v) := c(v) - k \cdot e(v)$.

Theorem 5.3. Let V be a finite set. Let $c : V \rightarrow \mathbb{N}$ and $k \geq 1$. If $|c| - (k - 1)(s^\#(c) - 1) \geq nk$, then n blocks of k can be extracted from c .

Proof. Use induction on n . The proof is trivial for $n = 0$. Suppose the theorem holds for n , and $|c| - (k - 1)(s^\#(c) - 1) \geq (n + 1) \cdot k$. Note that $|c| > s^\#(c) \cdot (k - 1)$. Using the pigeonhole principle determine $v \in V$ with $c(v) \geq k$. Subtract k from $c(v)$ and call the resulting configuration c^* . Now $|c^*| - (k - 1)(s^\#(c^*) - 1) \geq nk$, and using the induction hypothesis we can extract n blocks of k from c^* . In total $n + 1$ blocks of k are extracted from c , proving the theorem. \square

Corollary 5.4. Let V be a finite set. Let $c : V \rightarrow \mathbb{N}$ and $k \in \mathbb{N}$. If $n = \lceil (|c| - s^\#(c) \cdot (k - 1)) / k \rceil$ then n blocks of k can be extracted from c .

Proof. Apply the previous theorem. Note that $\lceil x/k \rceil = \lfloor (x + k - 1)/k \rfloor$ for any x . \square

In Theorem 5.3 you can think of $(k - 1)(s^\#(c) - 1)$ as a kind of offset that accounts for vertices not having an exact multiple of k pebbles. Any combination of blocks of at most k pebbles that can be fitted into the configuration size that remains after subtracting this offset can be extracted, no matter how spread out the configuration is. If all pebbles are placed on a single vertex, such that $s^\#(c) = 1$, then no offset is needed. To use this concept, we define the k -reduced size of a configuration:

Definition 5.5. The k -reduced size r_k of a configuration c is defined as:

$$r_k(c) := |c| - (k - 1)(s^\#(c) - 1)$$

Lemma 5.6. Let V be a finite set. For $c_1, c_2 : V \rightarrow \mathbb{N}$ holds $r_k(c_1 + c_2) \geq r_k(c_1) + r_k(c_2)$.

Proof. Left to the reader. Note that the support of c_1 and c_2 may intersect. \square

Lemma 5.7. Let V be a finite set. Let $c : V \rightarrow \mathbb{N}$ and $k, n \in \mathbb{N}$. If $r_k(c) \geq n$ and $n \leq k$, then one block of n pebbles can be extracted from c , resulting in c' with $r_k(c') \geq r_k(c) - n$.

Proof. Left to the reader. Note that the support may decrease. \square

5.2 The Induction Step

Our alternative to property (ii) of Theorem 5.1 uses the concept of reduced configuration size. Instead of directly extracting pebbles, we prove that two pebbles can be moved to a target vertex in such a way that the configuration which is left over has a certain k -reduced size. This will ensure that enough pebbles can be extracted from it later. We define this property in the form of a new kind of pebbling number written as $\tau_{n,k}$.

Definition 5.8. Let G be a graph. Given configurations c_1 and c_2 on G , we say that c_1 is a subconfiguration of c_2 , denoted as $c_1 \subseteq c_2$, if $c_1(v) \leq c_2(v)$ for all $v \in V(G)$.

Definition 5.9. Let G be a graph. The number $\tau_{n,k}(G, t)$ is the smallest $p \in \mathbb{N}$ where for every configuration c with $|c| = p - s^\#(c) + 1 + m$, there is a subconfiguration c^* of c such that c^* is n -fold t -solvable, and the k -reduced size of the residual configuration $c - c^*$ is at least m :

$$\tau_{n,k}(G, t) := \min \left\{ p \in \mathbb{N} \mid \begin{array}{l} \forall m \in \mathbb{N} \forall c : |c| = p - s^\#(c) + 1 + m \Rightarrow \\ \exists c^* \subseteq c : \text{Solv}_{n,t}(c^*) \wedge r_k(c - c^*) \geq m \end{array} \right\}$$

Example Let's illustrate the definition of $\tau_{n,k}$ on the graph $G = P_2^{(3)} \square P_2^{(4)}$. We will denote configurations on G using a square matrix, where steps in the horizontal direction have weight 4 and steps in the vertical direction have weight 3. The target vertex t will be the lower left corner. Later in this chapter it will be easy to see that the following holds:

$$\tau_{2,3}(P_2^{(3)} \square P_2^{(4)}, t) \leq 2 \cdot 3 \cdot 4 = 24$$

According to this bound, every configuration c on G with $|c| = 24 - s^\#(c) + 1 + m$ has a 2-fold t -solvable subconfiguration such that the residual configuration has a 3-reduced size of at least m . If we are given a configuration that puts 39 pebbles in the top right corner, so that $m = 15$, then we must pick the subconfiguration with 24 pebbles on this corner. This subconfiguration is 2-fold t -solvable, and the residual configuration has a 3-reduced size of 15 (equal to its normal size because the support is one), which is sufficient.

$$r_3 \left(\begin{bmatrix} 0 & 39 \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} 0 & 24 \\ 0 & 0 \end{bmatrix} \right) = 15 \geq m = 15$$

If we are instead given a configuration that is more spread out, say 13 pebbles on each of the corners except t , then we have more freedom. Note that now $m = 17$, since the support of this configuration contains 3 vertices. We could for example use one of the corners to form a subconfiguration that is 2-fold t -solvable, and still have plenty of pebbles left to make sure that the residual configuration has a sufficient 3-reduced size.

$$r_3 \left(\begin{bmatrix} 13 & 13 \\ 0 & 13 \end{bmatrix} - \begin{bmatrix} 13 & 0 \\ 0 & 0 \end{bmatrix} \right) = 24 \geq m = 17$$

5 Hypercubes

Lemma 5.10. *If v is the only vertex of K_1 , then:*

$$\tau_{2,k}(K_1, v) \leq 2$$

Proof. Let c be a configuration on K_1 with $|c| = 2 - s^\#(c) + 1 + m = 2 + m$. Define the subconfiguration c^* as $c^*(v) = 2$. Now c^* is 2-fold v -solvable and $r_k(c - c^*) = m$. \square

Lemma 5.11. *Let G be a graph. If $\pi(G, t) \leq p$ and $\tau_{2,k}(G, t) \leq 2p$, then for $n \geq 2$:*

$$\tau_{n,k}(G, t) \leq np$$

Proof. Let c be a configuration on G with $|c| = np - s^\#(c) + 1 + m$. First move two pebbles to t using $\tau_{2,k}(G, t) \leq 2p$, and then move another $n - 2$ pebbles to t from the residual configuration using $\pi(G, t) \leq p$. The details are left to the reader. \square

Our goal is to prove Theorem 5.1 for hypercubes with generalized edge weights. We prove the induction step as two separate lemmas. Instead of two n -cubes forming an $(n + 1)$ -cube, we will have two copies of an arbitrary graph with edges between matching vertices. To record the fact that between these two identical graphs, which were previously called M_1 and M_2 , pebbles are only moved in one direction, we define the *arrow graph*; a single directed edge with weight k going from 0 to 1.

Definition 5.12. *Let $k \in \mathbb{N}$. The arrow graph denoted by $\downarrow^{(k)}$ is defined as (V, E, ω) where:*

$$V = \{0, 1\}, \quad E = \{(0, 1)\}, \quad \omega(0, 1) = k.$$




Figure 20 – The Cartesian product of 3 arrow graphs.




Figure 21

Lemma 5.13. *Let G be a graph and t a target vertex. Let $k, l \in \mathbb{N}$ such that $2 \leq k \leq l$. If $\pi(G, t) \leq p$ and $\tau_{2,l}(G, t) \leq 2p$ then:*

$$\pi(\downarrow^{(k)} \square G, (1, t)) \leq kp$$

Proof. Let c be a configuration on $\downarrow^{(k)} \square G$ such that $|c| \geq kp$. Split the graph into G_0 containing the vertex $(0, t)$ and G_1 containing the vertex $(1, t)$. Split c into the configurations c_0 on G_0 and c_1 on G_1 . Consider the following two cases:

(i) $|c_1| \geq s^\#(c_0)$.

- Use Lemma 5.7 to move $\lfloor r_k(c_0)/k \rfloor$ pebbles to G_1 . Note that $k \leq l$.
- Derive $|c_1| + \lfloor r_k(c_0)/k \rfloor \geq p$ from the following inequalities:

$$\begin{aligned} r_k(c_0) &\leq k \lfloor r_k(c_0)/k \rfloor + (k-1) \\ k|c_1| + r_k(c_0) &= k|c_1| + |c_0| - (k-1)(s^\#(c_0) - 1) \\ &\geq k|c_1| - (k-1)|c_1| + (k-1) + |c_0| \\ &\geq kp + (k-1) \end{aligned}$$

- Now G_1 contains at least p pebbles, which is enough to move one pebble to $(1, t)$.

(ii) $|c_1| \leq s^\#(c_0) - 1$.

- Derive $|c_0| \geq kp - |c_1| \geq kp - s^\#(c_0) + 1$.
- Use Lemma 5.11 to move k pebbles to $(0, t)$, and from there move one pebble to $(1, t)$.

□

5 Hypercubes

Lemma 5.14. *Let G be a graph and t a target vertex. Let $k, l \in \mathbb{N}$ such that $2 \leq k \leq l$. If $\pi(G, t) \leq p$ and $\tau_{2,l}(G, t) \leq 2p$ then:*

$$\tau_{2,k}(\downarrow^{(k)} \square G, (1, t)) \leq 2kp$$

Proof. Let c be a configuration on $\downarrow^{(k)} \square G$ such that $|c| = 2kp - s^\#(c) + 1 + m$. Again split the graph into G_0 and G_1 , and split c into c_0 and c_1 . Consider the following two cases:

(i) $|c_1| \leq 2p - s^\#(c_1)$.

- Define $m^* := kp - |c_1| - s^\#(c_1)$.
- Derive $|c_0| \geq kp - s^\#(c_0) + 1$:

$$\begin{aligned} |c_0| &\geq 2kp - s^\#(c) + 1 - |c_1| \\ &= (kp - s^\#(c_0) + 1) + (kp - s^\#(c_1)) - |c_1| \\ &\geq kp - s^\#(c_0) + 1 \end{aligned}$$

- Move k pebbles to $(0, t)$ in G_0 using Lemma 5.11, resulting in a residual configuration c_0^* with $r_k(c_0^*) \geq m^* + m$:

$$\begin{aligned} r_k(c_0^*) &\geq |c_0| - kp + s^\#(c_0) - 1 \\ &= kp + |c_0| + s^\#(c_0) - 2kp - 1 \\ &= kp + |c| - |c_1| + s^\#(c) - s^\#(c_1) - 2kp - 1 \\ &= m^* + m \end{aligned}$$

- Move $\lfloor m^*/k \rfloor$ pebbles to G_1 using c_0^* and derive $|c_1| + \lfloor m^*/k \rfloor \geq p$. What remains of c_0^* has a k -reduced size of at least m .
- There are now at least p pebbles on G_1 , so one pebble can be moved to $(1, t)$ in G_1 , and another pebble can be moved to $(1, t)$ using the k pebbles that were moved to $(0, t)$.

(ii) $|c_1| \geq 2p - s^\#(c_1) + 1$.

- Move two pebbles to $(1, t)$ in G_1 using the hypothesis $\tau_{2,l}(G, t) \leq 2p$, resulting in a residual configuration c_1^* .
- Derive $r_k(c_0 + c_1^*) \geq m$:

$$\begin{aligned} s^\#(c_0) &\leq \#V(G) \leq p \\ s^\#(c_0) &\leq (k-1)(2p - s^\#(c_0) + 1) - 1 \\ r_k(c_0 + c_1^*) &\geq r_k(c_0) + r_k(c_1^*) \\ &\geq |c_0| - (k-1)(s^\#(c_0) - 1) + |c_1| - 2p + s^\#(c_1) - 1 \\ &= |c| - 2kp + s^\#(c_1) + (k-1)(2p - s^\#(c_0) + 1) - 1 \\ &\geq |c| - 2kp + s^\#(c) = m \end{aligned}$$

□

5.3 Results

Theorem 5.15. For $2 \leq k_1 \leq k_2 \leq \dots \leq k_n$:

$$\pi(\downarrow^{(k_1)} \square \downarrow^{(k_2)} \square \dots \square \downarrow^{(k_n)}, (1, 1, \dots, 1)) = k_1 k_2 \dots k_n$$

Proof. This theorem results from inductively applying Lemma 5.13 and 5.14. To prove the base case, note that $\downarrow^{(k_n)}$ is isomorphic to $\downarrow^{(k_n)} \square K_1$, and use Lemma 5.10. To see that the bound obtained so far is strict, put $k_1 k_2 \dots k_n - 1$ pebbles on $(0, 0, \dots, 0)$ and try to move one pebble to $(1, 1, \dots, 1)$. \square

Corollary 5.16. For $2 \leq k_1 \leq k_2 \leq \dots \leq k_n$:

$$\pi(P_2^{(k_1)} \square P_2^{(k_2)} \square \dots \square P_2^{(k_n)}) = k_1 k_2 \dots k_n$$

Proof. Use the previous theorem. \square

Generalization

Chung says that $\pi(K_n \square G) = n \cdot \pi(G)$ when all edge weights are 2 and G has the 2-pebbling property. We can state a similar generalization of the results we just presented using a star-like digraph with all edges pointing towards the central vertex. We denote this graph as $R_n^{(k)}$, where n is the number of leafs and every edge has weight k . Note that the arrow graph $\downarrow^{(k)}$ is equivalent to $R_1^{(k)}$. We label the central vertex of $R_n^{(k)}$ as s . The pebbling number of s can be determined using a proof similar to Theorem 2.10:

$$\pi(R_n^{(k)}, s) = nk - n + 1$$

Is the following also true?

Conjecture 1. Let G be a graph and $n \geq 1$. Let $k, l \in \mathbb{N}$ such that $2 \leq k \leq l$. If $\pi(G, t) \leq p$ and $\tau_{2,l}(G, t) \leq 2p$ then:

- (i) $\pi(R_n^{(k)} \square G, (s, t)) \leq (nk - n + 1) \cdot p$
- (ii) $\tau_{2,k}(R_n^{(k)} \square G, (s, t)) \leq (nk - n + 1) \cdot 2p$




Figure 22 – Inward pointed star graphs

6 Zero-sums

With the results that we obtained in the previous chapter, we are finally ready to prove the Erdős-Lemke conjecture (Theorem 1.1). To do this, we first show that pebbling solutions on one graph can be translated to solutions on another graph via a homomorphism. We use this approach to determine an upper bound for the pebbling number of a divisor lattice. Then we associate pebbles with subsequences, and use the pebbling steps to carry out a construction with these subsequences, which will finally produce the zero-sum subsequence we are looking for. Before starting, we note the following:

Theorem 6.1. *Every non-empty sequence of integers a_1, a_2, \dots, a_n has a non-empty subsequence b_1, b_2, \dots, b_m such that:*

$$b_1 + b_2 + \dots + b_m \equiv 0 \pmod{n}$$

Proof. Define $c_0 := 0$ and $c_i := (a_1 + a_2 + \dots + a_i) \pmod{n}$. Use the pigeonhole principle to determine i and j with $0 \leq i < j \leq n$ such that $c_i = c_j$. The subsequence $a_{i+1}, a_{i+2}, \dots, a_j$ has the desired property. \square

6.1 Homomorphisms

A homomorphism from a graph G to a graph H is a mapping ϕ from vertices of G to vertices of H that is *structure-preserving*, meaning that for every edge (u, v) in G there is also an edge $(\phi(u), \phi(v))$ in H . It is not necessary that H contains a copy of G ; sometimes it is possible to map multiple vertices in G to one vertex in H . For example, for all bipartite graphs there is a structure-preserving mapping to K_2 . We also want the homomorphism to preserve the edge weight, leading us to the following definition:

Definition 6.2. *Let G and H be graphs. A function $\phi : V(G) \rightarrow V(H)$ is a homomorphism from G to H , also written as $\phi : G \rightarrow H$, if for all $(u, v) \in E(G)$ we have:*

$$(\phi(u), \phi(v)) \in E(H) \quad \text{and} \quad \omega_G(u, v) = \omega_H(\phi(u), \phi(v)).$$

When a homomorphism $\phi : G \rightarrow H$ is surjective, we can use it to map a configuration c on H to a (not necessarily unique) configuration c^* on G in the following way: for every $v \in V(H)$, determine a vertex $u \in V(G)$ such that $\phi(u) = v$, and set $c^*(u) = c(v)$. Now a solution of c^* on G can be translated into a solution of c on H via the homomorphism.

Lemma 6.3. *Let G and H be graphs. If a surjective homomorphism $\phi : G \rightarrow H$ exists, then:*

- (i) $\pi(H) \leq \pi(G)$.
- (ii) $\pi(H, \phi(t)) \leq \pi(G, t)$ for all $t \in V(G)$.
- (iii) $\tau_{n,k}(H, \phi(t)) \leq \tau_{n,k}(G, t)$ for all $t \in V(G)$ and $n, k \in \mathbb{N}$.

Proof. Left to the reader. When proving (iii) it should be noted that the k -reduced size of a configuration on G does not decrease in the image of ϕ . □

Grids

In this section the vertices of a path $P_n^{(k)}$ are treated as numbers from 0 to $n - 1$. We denote the n -dimensional hypercube as Q_n and label its vertices with binary strings of zeros and ones. We define a special notation to count the number of ones in a binary string:

Definition 6.4. *Let x be a binary string. The number of ones in x is denoted by $\text{count}(x)$.*

Theorem 6.5. *For $2 \leq k_1 \leq k_2 \leq \dots \leq k_s$ and n_1, n_2, \dots, n_s with $n_i \geq 1$ for all $i \leq s$:*

$$\pi(P_{n_1}^{(k_1)} \square P_{n_2}^{(k_2)} \square \dots \square P_{n_s}^{(k_s)}) = k_1^{n_1-1} k_2^{n_2-1} \dots k_s^{n_s-1}$$

Proof. Define a homomorphism ϕ as follows:

$$\begin{aligned} \phi : Q_{n_1-1}^{(k_1)} \square Q_{n_2-1}^{(k_2)} \square \dots \square Q_{n_s-1}^{(k_s)} &\longrightarrow P_{n_1}^{(k_1)} \square P_{n_2}^{(k_2)} \square \dots \square P_{n_s}^{(k_s)} \\ \phi((x_1, x_2, \dots, x_s)) &:= (\text{count}(x_1), \text{count}(x_2), \dots, \text{count}(x_s)) \end{aligned}$$

Note that ϕ is a valid homomorphism, preserving both edge connectivity and weight, and that it is surjective. Using Lemma 6.3 and Corollary 5.16 we obtain an upper bound for the pebbling number of the product of paths. It is left to the reader to see that it is strict. □




Figure 23 – Retracting a hypercube to a 2×2 grid.

Divisor Lattices

The divisor lattice of a number n consists of all divisors of n as vertices, including n itself, and for each divisor an edge to every other divisor that is a multiple of it (excluding self-loops). In our definition, the edge weights are determined by the multiplication factor.

Definition 6.6. Let $n \in \mathbb{N}$. The divisor lattice D_n is a graph (V, E, ω) where:

$$\begin{aligned} V &= \{d \in \mathbb{N} \mid (d \mid n)\} \\ E &= \{(p, d) \in V \times V \mid p \neq d \wedge (p \mid d)\} \\ \omega(p, d) &= d/p \end{aligned}$$

We want to know the pebbling number of this graph. By looking at the prime factorization of n 's divisors, we can see that D_n contains a grid with the length of each dimension being the exponent of one of n 's prime factors plus one. Using the result from the previous section we can then determine that $\pi(D_n, n) \leq n$. This is essentially the approach taken by Chung. To make it easier to formalize, we instead prove this result directly from the step lemmas that we also used in the previous section.

Theorem 6.7. Let $n \geq 2$. If p is the smallest prime factor of n , then:

$$\pi(D_n, n) \leq n \quad \text{and} \quad \tau_{2,k}(D_n, n) \leq 2n.$$

Proof. We prove this theorem by induction on n . Suppose the theorem has been proven for all $m < n$. Let $m := n/p$, and define a homomorphism $\phi : \downarrow^{(p)} \square D_m \rightarrow D_n$ as follows:

$$\begin{aligned} \phi((0, d)) &:= d \\ \phi((1, d)) &:= pd \end{aligned}$$

Note that ϕ is well-defined and surjective. Using Lemma 6.3 we find:

$$\begin{aligned} \pi(D_n, n) &\leq \pi(\downarrow^{(p)} \square D_m, (1, m)) \\ \tau_{2,k}(D_n, n) &\leq \tau_{2,k}(\downarrow^{(p)} \square D_m, (1, m)) \end{aligned}$$

Now use Lemma 5.13 and 5.14, the induction steps for π and $\tau_{2,k}$ from the previous chapter. If $m \geq 2$ then the proof can be finished using the induction hypothesis. Note that $p \leq p_m$ where p_m is the smallest prime factor of m .

Otherwise, if $m = 1$, use that D_1 is isomorphic to K_1 . □

6.2 Pebbling Construction

Now we can finally show how graph pebbling is connected to number theory. Pebbles will carry sequences of numbers, and a predicate will determine if a sequence is *well-placed* on a certain vertex, such that a pebble carrying that sequence may be put there. A pebbling step will represent a way to pick a subsequence from the union of a number of sequences. Moving one pebble to a target vertex t will result in the construction of a subsequence with the property that it is well-placed on t .

Lemma 6.8. *Let G be a graph and A an set. Let $a_1, a_2, \dots, a_n \in A$ be a sequence of elements in A . Define $I := \{1, 2, \dots, n\}$. For every $v \in V(G)$, let a predicate \mathcal{P}_v on subsets of I be given that defines if a subsequence $\{a_i\}_{i \in S}$, determined by a set $S \subseteq I$, is well-placed on v . Suppose:*

- (I) *For every $i \in I$ there is a vertex v such that $\mathcal{P}_v(\{i\})$.*
- (II) *For every edge (u, v) in G and every sequence S_1, S_2, \dots, S_k of disjoint and non-empty subsets of I , where $k = \omega_G(u, v)$ and $\mathcal{P}_u(S_i)$ for all $i \leq k$, there exists a non-empty set $S \subseteq S_1 \cup S_2 \cup \dots \cup S_k$ such that $\mathcal{P}_v(S)$.*

If $\pi(G, t) \leq n$ for a vertex t , then there exists a non-empty set $S \subseteq I$ such that $\mathcal{P}_t(S)$.

Proof. Put every a_i on a vertex determined by (I). Define an initial pebbling configuration c by counting how many a_i 's are placed on each vertex, and note that c is t -solvable since $|c| = n$. Using (II) and a solution of c that puts one pebble on t , we construct a subsequence of a_1, a_2, \dots, a_n that is well-placed on t . \square

Divisor Pebbles

To demonstrate the concept of a pebbling construction more clearly, we are first going to prove the Erdős-Lemke conjecture for $n = d$:

Theorem 6.9. *Let $n \geq 1$. Every sequence a_1, a_2, \dots, a_n of n divisors of n has a non-empty subsequence b_1, b_2, \dots, b_m such that:*

$$b_1 + b_2 + \dots + b_m = n$$

Proof. Use Lemma 6.8 with $G = D_n$ and the sequence a_1, a_2, \dots, a_n . Let a subsequence be well-placed on a vertex $d \in V(D_n)$ if its sum is equal to d . For each $i \leq n$ we can place the one-element sequence a_i on the vertex a_i . If a vertex $d \in V(D_n)$ has at least p pebbles, in the form of p disjoint subsequences that each sum to d , then the union of these subsequences sums to pd , allowing us to put one pebble on the vertex pd in exchange for removing p pebbles from vertex d . This proves (I) and (II) of Lemma 6.8. Since $\pi(D_n, n) \leq n$, it follows that there is a subsequence of a_1, a_2, \dots, a_n which sums to n . \square

Example We consider the case $n = 60 = 2^2 \cdot 3 \cdot 5$. The configuration shown below, on the Hasse diagram of D_{60} , is one possible way of picking 60 divisors of 60: 29 ones, 15 twos, 10 threes, and 6 fives. The edge weight is 2 for going left, 3 for going up, and 5 for going right. A solution of this configuration, for target vertex 60, is shown as a pebble flow. We put 5 pebbles on 6 using 15 twos, and another 5 using 10 threes, then we use the 10 pebbles on 6 to put 2 pebbles on 30, and finally we can put one pebble on 60.




Figure 24

GCD Pebbles

To prove the case $n > d$, Lemke and Kleitman used a generalization of the theorem based on *greatest common divisors*. We will prove this generalization following the approach of Chung, including a technical correction described in [6]. We made a computer formalization of this proof using the *Coq Proof Assistant*, following the same steps as described in this and the previous chapter. The formalization required investigating all reasoning steps in detail, and gives an almost absolute guarantee that the proof, which we are about to finish, is correct. More details about the formalization can be found in the appendix.

Theorem 6.10. *Let $n \geq 1$. For every sequence of positive integers a_1, a_2, \dots, a_n there is a non-empty subset S of $\{1, 2, \dots, n\}$ such that:*

$$\sum_{i \in S} a_i \equiv 0 \pmod{n} \quad \text{and} \quad \sum_{i \in S} \gcd(n, a_i) \leq n.$$

Proof. We again use Lemma 6.8. This time, a subsequence determined by $S \subseteq \{1, 2, \dots, n\}$ is well-placed at vertex $d \in V(D_n)$ if:

- (a) $\sum_{i \in S} a_i$ is a multiple of d ,
- (b) $\sum_{i \in S} \gcd(n, a_i) \leq d$.

It is now sufficient to prove property (I) and (II) of Lemma 6.8:

- (I) For all $i \leq n$ the one-element sequence a_i is well-placed at vertex $\gcd(n, a_i)$.
- (II) Let (d, pd) be an edge of D_n . Note that $\omega(d, pd) = p$. Suppose a sequence of p disjoint sets $S_1, S_2, \dots, S_p \subseteq \{1, 2, \dots, n\}$ is given that are all well-placed at vertex d . Define:

$$\begin{aligned} x_i &:= \sum_{j \in S_i} a_j \\ r_i &:= x_i/d \end{aligned}$$

Note that x_i is a multiple of d according to (a). Using Theorem 6.1, determine a set $R \subseteq \{1, 2, \dots, p\}$ corresponding to a subsequence of r_1, r_2, \dots, r_p that sums to a multiple of p . Now the set S given by $\bigcup_{i \in R} S_i$ is well-placed at vertex pd .

□

6.2 Pebbling Construction

Corollary 6.11. (Erdős-Lemke conjecture) *Given positive integers n , d , and a_1, a_2, \dots, a_d with $d \mid n$ and $a_i \mid n$ for all $i \leq d$, there is a non-empty subset S of $\{1, 2, \dots, d\}$ such that:*

$$\sum_{i \in S} a_i \equiv 0 \pmod{d} \quad \text{and} \quad \sum_{i \in S} a_i \leq n.$$

Proof. Apply the previous theorem with $n := d$. For the second requirement, use:

$$\sum_{i \in S} a_i = \sum_{i \in S} \gcd(n, a_i) \leq \sum_{i \in S} n/d \cdot \gcd(d, a_i) \leq n$$

This is based on the following property of gcd with $k = n/d$:

$$\gcd(k \cdot d, a) \leq k \cdot \gcd(d, a)$$

□





Figure 25 – Paul Erdős, Ronald Graham and Fan Chung (1986).

Bibliography

- [1] D. H. Mellor. “The Eponymous F. P. Ramsey”. In: *Journal of Graph Theory* 7.1 (1983), pp. 9–13. doi: 10.1002/jgt.3190070103.
- [2] Paul Lemke and Daniel Kleitman. “An addition theorem on the integers modulo n ”. In: *Journal of Number Theory* 31.3 (1989), pp. 335–345. issn: 0022-314X. doi: 10.1016/0022-314X(89)90077-2.
- [3] Fan R. K. Chung. “Pebbling in Hypercubes”. In: *SIAM Journal on Discrete Mathematics* 2.4 (1989), pp. 467–472. doi: 10.1137/0402041.
- [4] David Moews. “Pebbling graphs”. In: *Journal of Combinatorial Theory, Series B* 55.2 (1992), pp. 244–252. issn: 0095-8956. doi: 10.1016/0095-8956(92)90043-W.
- [5] B. Voxman L. Pachter H. Snevily. “On pebbling graphs”. In: *Congressus Numerantium* 107 (1995), pp. 65–80. url: <https://resolver.caltech.edu/CaltechAUTHORS:20170309-150137723>.
- [6] T. A. Clarke, R. A. Hochberg, and G. H. Hurlbert. “Pebbling in Diameter Two Graphs and Products of Paths”. In: *Journal of Graph Theory* 25.2 (1997), pp. 119–128. issn: 0364-9024. url: <https://dl.acm.org/doi/10.5555/252979.252986>.
- [7] David S. Herscovici and Aparna W. Higgins. “The Pebbling Number of $C_5 \times C_5$ ”. In: *Discrete Mathematics* 187.1–3 (1998), pp. 123–135. issn: 0012-365X. doi: 10.1016/S0012-365X(97)00229-X.
- [8] Hunter S. Snevily and James A. Foster. “The 2-Pebbling Property and a Conjecture of Graham’s”. In: *Graphs and Combinatorics* 16.2 (2000), pp. 231–244. issn: 1435-5914. doi: 10.1007/PL00021179.
- [9] Kevin Milans and Bryan Clark. “The Complexity of Graph Pebbling”. In: *SIAM Journal on Discrete Mathematics* 20.3 (2006), pp. 769–798. doi: 10.1137/050636218.
- [10] David S. Herscovici. “Graham’s Pebbling Conjecture on Products of Many Cycles”. In: *Discrete Mathematics* 308.24 (2008), pp. 6501–6512. issn: 0012-365X. doi: 10.1016/j.disc.2007.12.045.

Bibliography

- [11] Nándor Sieben. “A Graph Pebbling Algorithm on Weighted Graphs”. In: *Journal of Graph Algorithms and Applications* 14.2 (2010), pp. 221–244. DOI: 10.7155/jgaa.00205.
- [12] Glenn Hurlbert. *A linear optimization technique for graph pebbling*. 2011. DOI: 10.48550/arxiv.1101.5641.
- [13] Roberto Cavada *et al.* *NuSMV 2.6 User Manual*. FBK-irst. 2015. URL: <https://nusmv.fbk.eu/NuSMV/userman/v26/nusmv.pdf>.
- [14] Daniel W. Cranston *et al.* “Modified Linear Programming and Class 0 Bounds for Graph Pebbling”. In: *Journal of Combinatorial Optimization* 34.1 (2017), pp. 114–132. ISSN: 1382-6905. DOI: 10.1007/s10878-016-0060-6.
- [15] Glenn Hurlbert. “The Weight Function Lemma for Graph Pebbling”. In: *Journal of Combinatorial Optimization* 34.2 (2017), pp. 343–361. ISSN: 1382-6905. DOI: 10.1007/s10878-016-9993-z.
- [16] Ze-Tu Gao and Jian-Hua Yin. “Lemke graphs and Graham’s pebbling conjecture”. In: *Discrete Mathematics* 340.9 (2017), pp. 2318–2332. ISSN: 0012-365X. DOI: 10.1016/j.disc.2016.09.004.
- [17] Ervin Györi, Gyula Y. Katona, and László F. Papp. “Constructions for the Optimal Pebbling of Grids”. In: *Periodica Polytechnica Electrical Engineering and Computer Science* 61.2 (2017), pp. 217–223. DOI: 10.3311/PPee.9724.
- [18] Ervin Györi, Gyula Y. Katona, and László F. Papp. “Optimal Pebbling Number of the Square Grid”. In: *Graphs and Combinatorics* 36.3 (2020), pp. 803–829. ISSN: 1435-5914. DOI: 10.1007/s00373-020-02154-z.
- [19] Franklin Kenter, Daphne Skipper, and Dan Wilson. “Computing bounds on product graph pebbling numbers”. In: *Theoretical Computer Science* 803 (2020), pp. 160–177. ISSN: 0304-3975. DOI: 10.1016/j.tcs.2019.09.050.
- [20] Glenn Hurlbert and Franklin Kenter. “Graph Pebbling: A Blend of Graph Theory, Number Theory, and Optimization”. In: *Notices of the American Mathematical Society* 68.11 (2021). DOI: 10.1090/noti2379.

Appendix: Coq Formalization

The purpose of this appendix is to show some of the highlights of our Coqⁱ formalization of the Erdős-Lemke conjecture. The source code of this formalization is available onlineⁱⁱ. Here we use some slightly different notations than in the Coq source to improve readability. The formalization was developed using version 8.16.0 of Coq, and version 1.8.0 of the the *Coq-std++* libraryⁱⁱⁱ. Future versions of these systems may introduce breaking changes, and so it is unsure if our formalization can still be evaluated a few decades from now.

```
Corollary Erdos_Lemke_conjecture (l : list nat) (d n : nat) :
  length l = d → d > 0 → n > 0 → d | n → (∀ a, a ∈ l → a | n) →
  ∃ (l' : list nat), l' ≠ [] ∧ l' ⊆+ l ∧ d | summation l' ∧ summation l' ≤ n.
```

Figure 26 – The Erdős-Lemke conjecture in Coq.

Naming Heuristics Like in most programming languages, all definitions and theorems in Coq must have a unique identifier expressed as a plain text string with various restrictions. Finding a naming pattern for these identifiers that is both coherent and convenient is a significant challenge when developing a large computer formalization. One of the naming *heuristics* (it is too loose to be called a convention) that we used to form new identifiers is appending the main identifiers in the new definition or theorem in their order of appearance. We proved many small lemmas about summations, of which three simple examples are shown below. Here \subseteq^+ denotes the multiset inclusion operator from Coq-std++.

```
Lemma elem_of_summation : n ∈ l → n ≤ summation l.
Lemma submseteq_summation : l1 ⊆+ l2 → summation l1 ≤ summation l2.
Lemma summation_nonzero : summation l > 0 → ∃ n, n ∈ l ∧ n > 0.
```

Figure 27 – Lemmas about the summation function.

ⁱ<https://coq.inria.fr/>

ⁱⁱ<https://github.com/bergwerf/pebbling>, DOI: 10.5281/zenodo.7876086

ⁱⁱⁱ<https://gitlab.mpi-sws.org/iris/stdpp>

Pebbling Bounds In the formalization it is not convenient to define the pebbling number as a function. We instead use a predicate that defines when, for a given target vertex, the pebbling number is *bounded* by a certain value:

- $\pi(G, t) \leq p$ is defined by: `vertex_pebbling_bound G p t`,
- $\tau_{n,k}(G, t) \leq p$ is defined by: `vertex_pebbling_property G n k p t`.

Using these predicates, we can formulate the main lemmas from Chapter 5 in Coq as shown in the excerpt below. Here the graph called ‘arrow’ is analogous to the arrow graph $\downarrow^{(k)}$.

Hypothesis $H_1 : 2 \leq k \leq l$.

Hypothesis $H_2 : \text{vertex_pebbling_bound } G \ p \ t$.

Hypothesis $H_3 : \text{vertex_pebbling_property } G \ 2 \ l \ (2 \cdot p) \ t$.

Lemma `vertex_pebbling_bound_arrow_prod` :
`vertex_pebbling_bound (arrow □ G) (k · p) (v1, t)`

Lemma `vertex_pebbling_property_arrow_prod` :
`vertex_pebbling_property (arrow □ G) 2 k (2 · k · p) (v1, t)`

Figure 28 – Results analogous to Lemma 5.13 and 5.14.

Homomorphisms Transferring pebbling bounds between graphs using surjective homomorphisms was surprisingly difficult to formalize. The main difficulty is the inverse mapping of configurations that we described after Definition 6.2.

Hypothesis `surj` : `Surj (=) h`.

Hypothesis `hom` : `Graph_Hom G H h`.

Lemma `surj_hom_vertex_pebbling_bound` $n \ t$:
`vertex_pebbling_bound G n t → vertex_pebbling_bound H n (h t)`.

Figure 29 – A result similar to Lemma 6.3(ii).

In Coq, we defined a small functional program called ‘inv_dmap’ (*inverse domain mapping*) that carries out this inversion procedure. Among other properties, we had to prove that the size and the number of elements in the support remain unchanged when this function is applied to a configuration.

Definition $\text{inv} (f : A \rightarrow B) (b : B) : \text{list } A :=$
 $\text{filter } (\lambda a, f a = b) (\text{enum } A).$

Definition $\text{dmap} (f : A \rightarrow \text{nat}) : B \rightarrow \text{nat} :=$
 $\lambda b, \text{summation } (f \langle \$ \rangle \text{inv } h b).$

Definition $\text{inv_dmap} (f : B \rightarrow \text{nat}) : A \rightarrow \text{nat} :=$
 $\lambda v, \text{if decide } (\text{head } (\text{inv } h (h v)) = \text{Some } v) \text{ then } f (h v) \text{ else } 0.$

Figure 30 – Configuration mappings.

Pebbling Construction The construction described in Lemma 6.8 is carried out in Coq using lists of elements of an arbitrary type A . A list called ‘pebbles’ determines the values that may be used. Each vertex in the graph is associated with any number of lists with elements of type A , and the concatenation of all these lists should be a permutation of (a part of) the pebbles list.

Variable $\text{pebbles} : \text{list } A.$

Variable $\text{well_placed} : V G \rightarrow \text{list } A \rightarrow \text{Prop}.$

Hypothesis $\text{initialize} : \forall a, a \in \text{pebbles} \rightarrow \exists v, \text{well_placed } v [a].$

Hypothesis $\text{step} : \forall (u v : V G) (l^* : \text{list } (\text{list } A)),$
 $E G u v \rightarrow \text{Forall } (\text{well_placed } u) l^* \rightarrow \text{length } l^* = \text{weight } G u v \rightarrow$
 $\exists (l : \text{list } A), l \subseteq^+ \text{concat } l^* \wedge \text{well_placed } v l.$

Lemma $\text{pebbling_construction } t :$
 $\text{vertex_pebbling_bound } G (\text{length } \text{pebbles}) t \rightarrow$
 $\exists (l : \text{list } A), l \subseteq^+ \text{pebbles} \wedge \text{well_placed } t l.$

Figure 31 – A result analogous to Lemma 6.8.