

Geometry Of The Subset Sum Problem - Part I

Srinivas Balaji Bollepalli

November 21, 2025

Abstract

We announce two breakthrough results concerning important questions in the Theory of Computational Complexity. In this expository paper, a systematic and comprehensive geometric characterization of the Subset Sum Problem is presented. We show the existence of a universal geometric structure, comprised of a family of non-decreasing paths in the Cartesian plane, that captures any instance of the problem of size n . Inspired by the geometric structure, we provide an unconditional, deterministic and polynomial time algorithm, albeit with fairly high complexity, thereby showing that $\mathcal{P} = \mathcal{NP}$. Furthermore, our algorithm also outputs the number of solutions to the problem in polynomial time, thus leading to $\mathcal{FP} = \#\mathcal{P}$. As a bonus, one important consequence of our results, out of many, is that the quantum-polynomial class $\mathcal{BQP} \subseteq \mathcal{P}$.

Not only this, but we show that when multiple solutions exist, they can be placed in certain equivalence classes based on geometric attributes, and be compactly represented by a polynomial sized directed acyclic graph. We show that the Subset Sum Problem has two aspects, namely a *combinatorial* aspect and a *relational* aspect, and that it is the latter which is the primary determiner of complexity. We reveal a surprising connection between the size of the elements and their number, and the precise way in which they affect the complexity. In particular, we show that for all instances of the Subset Sum Problem, the complexity is independent of the size of elements, once the difference between consecutive elements exceeds $\lceil 7 \log n \rceil$ bits in size.

We derive a metric to measure the degree of additive disorder (equivalently additive structure) in a given instance, based on some geometric translations, and show how this metric determines the computational complexity. We present examples of instances with varying degrees of additive structure to show this connection explicitly. We also discuss a relation between the geometric structure and the additive structure of the Subset Sum Problem. Along with our analysis technique, these results may be of general interest to *Additive Combinatorics*.

We provide some numerical examples to illustrate the algorithm, and also show how it can be used to estimate some difficult combinatorial quantities such as the number of restricted partitions.

Contents

1	Introduction	4
1.1	Previous approaches	4
1.2	An important question	5
1.3	About this work & presentation style	6
1.4	Ideas & Outline of the paper	7
2	Problem statement and notation	8
2.1	A running example	9
2.2	Subspaces	10
2.3	Additional Definitions	10
3	Filling boxes, and paths from 0 to 1	11
3.1	A model for filling n boxes	12
3.2	Two complementary ways of filling boxes	13
3.3	Formulae to generate NDS	14
3.4	From sequences to curves	17
3.5	Reflection Symmetry	17
3.6	The curves p_k and q_k and coverage	18
4	Links, Chains and Transforms	19
4.1	Elemental representation of chains	22
4.2	Transforms on chains	23
4.3	$p_k \rightarrow \{q_i\}_{i=1}^k$ and $q_k \rightarrow \{p_i\}_{i=1}^k$ structures	26
4.4	Nested τ operations	26
5	Generation of all NDPs	28
5.1	A transformation graph $\mathcal{H}$	28
5.2	Counting paths in $\mathcal{H}$	29
5.3	A covering by NDPs	34
5.4	Segments and Path multiplicity	36
5.5	A transformation vector	36
5.6	Wormhole paths of $\mathcal{H}$	38
5.7	Number of OLs into each curve/level for a given $T \in p_n$	40
6	Intersection Points and Local Coordinates	41
6.1	Edges and Local Reference Frames	41
6.2	Local coordinates of intersection points	42
6.3	A Bit Of Relativity	43
6.4	Interacting edges	44
6.5	An unique representation of an edge	45
6.6	Number of interacting edge pairs	46

7	The Orbital Graph G_0	47
7.1	Determining the root node e_0	48
7.2	The Graph G_0	49
7.3	Computation of arcs between complementary curves	49
7.4	Taking character sets into account: appending adjacency lists	50
7.5	Computation of nodes	50
7.6	Size of G_0	51
7.7	Paths in G_0	52
7.8	Classification of paths through a node	53
8	Hopping On The Orbital Line: An Algorithm	54
8.1	Single Source Shortest and Longest Paths	54
8.2	Refining Nodes and Arcs	54
8.3	A destination node e_∞	57
8.4	Filtering of nodes and arcs	58
8.5	The algorithm at a glance	61
8.6	Growth factor	62
8.7	Indices as path lengths	63
9	A Holy Trinity: Number, Size and Additive Structure	64
9.1	A Configuration Graph	64
9.2	Apparent paths in $\mathcal{C}$	67
9.3	Path magnification in $\mathcal{C}$	67
9.4	Maximum number of paths in $\mathcal{C}$	68
9.5	Some simple relations on γ_r and $ C(r, r + 1) $	69
9.6	The importance of size	69
9.7	A Product Inequality	70
9.8	An upper bound on $\max(\gamma)$	73
9.9	Profile of γ	75
9.10	Asymptotic profile of γ	76
9.11	Geometric Interpretation	79
9.12	The case of Zero Paths	79
10	Main Results & Complexity Calculations	83
10.1	Finding the number of solutions	84
10.2	Classifying the solutions	85
10.3	Numerical results	85
11	Instances & Additive Structure	85
11.1	Two quantities of interest	86
11.2	Dependence of Solution Graph size on U	87
11.3	A Special Case: $\mathbf{a} = \mathbf{b}$	88
11.4	A snapshot of instances	89

12 Epilogue	90
12.1 A combinatorial facade and a reconciliation	90
12.2 Historical significance	91
13 Appendix	95
13.1 Example 1: Multiple Solutions	95
13.2 Example 2: A Single Solution	96
13.3 Example 3: Growth of k_{peak} with n	98
13.4 Example 4: Varying m	98
13.5 Example 5: High degree of additive structure	100
13.6 Some Combinatorial Applications	100
13.6.1 Binomial Coefficients	100
13.6.2 Restricted Partitions I	101
13.6.3 Restricted Partitions II	101
13.7 A remark on size	101

1 Introduction

Given a sequence of positive numbers $(a_1, a_2, \dots, a_n)$, is there a subset that sums to a given target number T ? This is the famous *Subset Sum Problem (SSP)* that is easily understood by anyone who knows how to add two numbers. In fact, young school children not only understand it, but even successfully solve it when it involves just a handful of small numbers. All they have to do is try out various possible combinations, and check if there is a match, i.e., perform an *exhaustive search*. However, as n increases and the size of numbers gets bigger, the problem becomes hard quickly, and the famousness turns to notoriety, particularly when one fails to recognize any pattern. Could there be a pattern, in general, in the first place that obviates the need for an exhaustive search? We show that the rich geometric structure of the *Subset Sum Problem*, indeed helps us to avoid exponential search.

1.1 Previous approaches

There has been only an incremental improvement to the *exhaustive search* coming from the work of Horowitz and Sahni [15], with a time complexity of $O(n2^{\frac{n}{2}})$. It is a *divide-and-conquer* type of algorithm and consists of: (1) breaking the problem into (roughly) two halves, (2) enumerating all the subset sums for each half in increasing order and (3) checking for a combination of elements one from each half such that their sum is T . In this so called *two-list* algorithm, the computational effort of enumeration and sorting is $O(n2^{\frac{n}{2}})$ while checking for solution can be done in $O(2^{\frac{n}{2}})$. Therefore the overall complexity is $O(n2^{\frac{n}{2}})$. It is noteworthy that no better algorithm in terms of *worst-case* time complexity than this has been found since 1974. An improvement in terms of space was made in [23], where the worst case space needs were shown to be $O(n2^{\frac{n}{4}})$, but the time complexity remained at $O(n2^{n/2})$.

Following this idea of splitting the given set into two parts, and with the goal of beating the $O(n2^{n/2})$ bound, Howgrave-Graham and Joux [16] obtained an improved bound of $\tilde{O}(2^{0.337n})$ (the symbol $\tilde{O}$ means that some polynomial factors are suppressed) by splitting into four parts for instances of *average case* complexity. This was followed by another improvement by Becker, Coron and Joux [3], who showed a complexity of $\tilde{O}(2^{0.291n})$. These approaches being probabilistic in nature, do not always find the solution. More importantly, for instances with no solution, they cannot certify that there is no solution.

A pseudo-polynomial time algorithm with a time complexity of $O(n2^m)$ (where m represents the maximum number of bits needed to represent the numbers), based on *dynamic programming* [11; 8], can easily be constructed. Clearly, this is useful only when m is very small. Lagarias, Odlyzko and others [18; 7] looked at this problem as a *Lattice* problem, where *Subset Sum Problem* is reduced to a *shortest vector problem* (SVP) in a lattice. These methods are conditioned on the availability of efficient algorithms to find the *shortest vector* in a lattice. The currently known polynomial time algorithms provide a *short vector* but do not guarantee a *shortest vector*. As a result the lattice based methods do not always guarantee a solution, even when one exists.

Apart from these approaches, there is no known unconditionally deterministic algorithm that always solves the *Subset Sum Problem*. In addition, there haven't been any known efforts to characterize the structure of the SSP. Given this, there is no exaggeration in saying that nothing really is known about the structure of the SSP.

1.2 An important question

One might take some solace in the modest progress, by noting that the *SSP* is $\mathcal{NP}$ -complete. In other words, it is the hardest problem in $\mathcal{NP}$; The class of problems where given a solution, the verification is easy (i.e., can be done in polynomial time), but there is no known efficient method to find a solution in the first place. The $\mathcal{NP}$ -completeness of the *Subset Sum Problem* (along with several other combinatorial problems) was first proven in a seminal paper by Karp [17], following a landmark result of Cook [6], who in 1971 formally defined the notion of *NP-completeness* for the first time, by showing that the Boolean satisfiability problem *SAT* is $\mathcal{NP}$ -complete. This idea of $\mathcal{NP}$ -completeness was also arrived at independently by Levin in 1973 (see for e.g., the translation [28]), who called it a “universal sequential search problem” [11]. It turns out that any $\mathcal{NP}$ -complete problem can be transformed into another, through the notion of *reducibility* [17]. Using this idea of *reducibility*, hundreds of problems were proved to be $\mathcal{NP}$ -complete, and collected in the form of a book by Garey and Johnson [11]. The list of $\mathcal{NP}$ -complete problems is growing with new problems being identified on an almost regular basis.

On the other hand, $\mathcal{P}$ is the class of all problems that can be solved in polynomial time, i.e., in time n^c where c is a positive constant.

The $\mathcal{P}$ vs. $\mathcal{NP}$ problem asks whether these two classes are in fact equal. The famous logician Kurt Gödel wondered about this very question (in a letter to John von Neumann), as far back as 1956, even before these classes were formally defined [13; 25].

The theory of computational complexity has rapidly grown since the 70s, with the

discovery of many new complexity classes (apart from the main $\mathcal{P}$ and $\mathcal{NP}$) and continues to grow. Out of these, one particular class, namely $\#\mathcal{P}$ (sharp-P) is also relevant to our work. It refers to the class of all problems, called *enumeration problems* where we are interested in determining the number of solutions to a given instance. The $\#\mathcal{P}$ class was defined by Valiant [30], who showed that computing the permanent of a matrix is a complete problem for this class. Unlike the class $\mathcal{NP}$, where given a solution the verification can be done in polynomial time, here in the class $\#\mathcal{P}$ given the *number of solutions* to a given instance, this fact can be verified efficiently. While the class $\mathcal{NP}$ deals with *decision problems* (with a *yes/no* answer), the class $\#\mathcal{P}$ deals with *function problems* (the number of satisfying solutions). Just as *SSP* is a *complete* problem for the class $\mathcal{NP}$, it turns out that the counting version of *SSP* is *complete* for the $\#\mathcal{P}$ class [22]. If we knew the number of subsets that sum to the target number, then clearly we know the answer to the decision problem. Thus the counting version feels more *harder* than decision version. Just as $\#\mathcal{P}$ is the function counterpart to $\mathcal{NP}$, the counterpart to $\mathcal{P}$ is $\mathcal{FP}$ which is the class of all functions that can be computed in polynomial time.

Like the $\mathcal{P}$ vs. $\mathcal{NP}$ question, the function version asks if $\mathcal{FP} = \#\mathcal{P}$? Both these questions are answered in the affirmative in this paper.

Apart from the above brief remarks on the main complexity classes of $\mathcal{P}$, $\mathcal{NP}$, $\mathcal{FP}$, and $\#\mathcal{P}$, we will not further delve on it here. We do make some remarks on the historical context of the $\mathcal{P}$ vs. $\mathcal{NP}$, and its importance towards the end of the paper. There are entire books devoted to it [see for e.g., 22; 24; 2; 21], which can be consulted for more details. The history behind the $\mathcal{P}$ vs. $\mathcal{NP}$ question, its importance to *Mathematics* and *Computer Science* as well as to many other branches of science have been beautifully described in many works [19; 20; 24; 31; 1; 5]. A popular and entertaining account of the importance of this problem is narrated in [10].

To show that these classes are equal, one has to simply provide a polynomial time algorithm to *any one* of the thousands of $\mathcal{NP}$ -complete problems. A majority of the researchers in complexity theory, strongly believe that these classes are unequal. This belief is articulated in the above cited references, as well as many online blogs. One commonly cited reason for this belief is the *failure* in finding an efficient algorithm for any of the numerous $\mathcal{NP}$ -complete problems over the past 50+ years. We discuss more about this in Section 12.

1.3 About this work & presentation style

This work started with the primary goal to understand the *structure* of the *Subset Sum Problem*, with the view that it will provide an understanding of the structure of *all* $\mathcal{NP}$ -complete problems. We chose to approach this by considering *geometry of the solution space*, and fortunately it led to the discovery of a fascinating structure. Having met the initial goal of capturing the solution space by an *universal geometric structure*¹, the idea to devise an algorithm came quite naturally. Some obstacles appeared while assessing the algorithm, but they ended up revealing an *important relation* between the size of

¹A curious reader may take a sneak peek at Figure 8 and 18, and return.

the elements in the *Subset Sum Problem* and their number, and how they determine the complexity.

Due to the simplicity of the problem statement, the *Subset Sum Problem* has a universal appeal from hobbyists to career mathematicians. This fact along with the importance attached to the $\mathcal{P}$ vs. $\mathcal{NP}$ question, as well as the original and novel ideas in this work, prompted me to write the paper in an expository style, with a linear progression of ideas, at a somewhat leisurely pace by including a lot of details and several examples. Due to this adopted style, some readers may find it tedious at some places, while others may see increased clarity. However, the final judgment belongs to the reader if we have met the goals for clarity and correctness of claims.

1.4 Ideas & Outline of the paper

The ideas used in the paper are elementary in nature consisting of simple geometric concepts, sequences, combinatorics, paths and path lengths in a graph, shortest and longest paths, and a creative combination of them. Several proofs in the paper make use of the principle of mathematical induction, while some others make use of combinatorics and counting in two ways.

The following is the brief outline of the rest of the paper: Given that our main goal is to understand the *structure* of the *Subset Sum Problem*, roughly the first half of the paper (Sections 2 to 5) is devoted to the construction of a *universal geometrical structure* comprised of a family of non-decreasing paths in the Cartesian plane that captures the complete solution space of any instance. We show two different aspects of each non-decreasing path - a sequence of points and a sequence of *links* called as a *chain* - and their equivalence. We show that the non-decreasing paths admit two symmetries, namely a *reflection symmetry* and a *local character symmetry*. We describe transforms which convert one non-decreasing path to another. Armed with these transforms, we provide a systematic way to generate a whole family of non-decreasing paths starting with a single path. In this geometric characterization we have a family of curves that start at the origin and have a common end point. The target sum T corresponds to a horizontal line (*orbital line*) that intersects all the curves.

In Section 6 we analyze the intersection points with the *orbital line*. For this we introduce *local reference frames* and *local coordinates*. This is followed by treating each curve as a sequence of linear segments (called *edges*) and provide relations that translate an intersection point in one edge to another edge. We make the *edge representation* a central part of the characterization of the *orbital line* and construct a directed acyclic graph G_0 (Section 7) with *edges* as nodes, and *translations* as arcs. We show that the SSP is equivalent to finding *zero paths* in G_0 . We estimate the number of nodes and number of arcs in G_0 , and show that it is polynomial sized.

We present an iterative algorithm in Section 8 where we characterize the complexity in terms of a certain *growth factor* η_{peak} . In order to bound η_{peak} we take a deeper dive into the SSP in Section 9 by showing that it has two distinct aspects: a *combinatorial aspect* and a *relational aspect*. We construct a *configuration graph* that captures these two aspects, and by a systematic analysis obtain a bound for the maximum number of

distinct paths through an edge. We also show that *size doesn't matter beyond a critical value* dependent on $\log n$. In Section 10 we prove the main claims of the paper. In Section 11 we consider a set of instances with varying *degrees of additive structure* and show how the algorithm performs on them. We also try to provide a connection between the *geometric* and *additive* structures. We finally conclude by discussing the historical antecedents of the problem and the significance of our results. We present several numerical examples along with applications of the SSP algorithm to computing some difficult combinatorial quantities such as the number of restricted partitions in Section 13.

2 Problem statement and notation

We use the standard notation of $\mathbb{N}, \mathbb{R}, \mathbb{C}$ denoting the positive integers, reals and complex numbers respectively. For a positive integer k , the set $\{1, 2, \dots, k\}$ is denoted by $[k]$. For an integer k lying in between integers i and j , where $i < j$, we write $k \in [i, j]$.

Given $(\mathbf{a}, T)$, where $\mathbf{a} = (a_1, a_2, \dots, a_n)$ is a sequence of n positive integers and T is another given positive integer, we are interested in finding the *number of subsets*, if any, that sum to T . In contrast to the problem statement in Section 1, which is a *decision problem*, we are interested in the *counting version* here. We assume, without loss of generality that

$$a_1 \leq a_2 \leq \dots \leq a_{n-1} \leq a_n,$$

and that $0 < T \leq \sum_j a_j$. We further assume that each of the numbers a_j can be represented using $m \in \mathbb{N}$ digits in binary, where m is independent of n . In other words, $a_n < 2^m$. Let $b_j := 2^{j-1}$, where $j \in [n]$, be the binary basis sequence $\mathbf{b} = (2^0, 2^1, \dots, 2^{n-1})$.

DEFINITION 2.1. Let $c_j := b_j + \iota a_j$, where $\iota = \sqrt{-1}$, for each $j \in [n]$ be the n complex numbers formed from the given set $\mathbf{a}$ and the binary basis $\mathbf{b}$. We denote this sequence of n complex numbers by $\mathbf{c}$, i.e.,

$$\mathbf{c} = (c_1, c_2, \dots, c_n).$$

We denote the real part by $\Re$ and imaginary part by $\Im$, so that $\Re(c_1) = b_1$ and $\Im(c_1) = a_1$.

DEFINITION 2.2. For a positive integer $k \in [n]$, we define

$$\begin{aligned} B_k &:= b_1 + b_2 + \dots + b_k = 2^k - 1, \\ A_k &:= a_1 + a_2 + \dots + a_k, \\ \text{and } C_k &:= B_k + \iota A_k. \end{aligned}$$

DEFINITION 2.3. (Power Set) Given $\mathbf{c}$, we denote the power set by S_n which is defined as the set of all subset sums of $\mathbf{c}$. Formally, we have

$$S_n := \{x_1 c_1 + x_2 c_2 + \dots + x_n c_n : \mathbf{x} \in \{0, 1\}^n\},$$

where $\mathbf{x}$ is a binary vector of length n . It is clear that the cardinality of S_n denoted by $|S_n| = 2^n$. These subset sums can be represented by 2^n points in the complex plane $\mathbb{C}$. We also refer to the power set sometimes as point set. Each point of the point set corresponds to a unique binary vector $\mathbf{x}$ of length n so that the index $= \mathbf{b} \cdot \mathbf{x}$ and the sum $= \mathbf{a} \cdot \mathbf{x}$ where $\cdot$ indicates dot product.

Since $\mathbb{C} \cong \mathbb{R}^2$, there should not be any confusion when we frequently treat the complex numbers as ordered pairs in $\mathbb{R}^2$. It should be emphasized that, we use the term subset sum for both the ordinary integer sum as well as the complex sum that has index as the real part, and ordinary subset sum as the imaginary part. It should be clear from the context, what is being referred to.

The power set S_n is our universe, in this article, for studying the *Subset Sum Problem*. For the given positive integer T , we can imagine drawing a line $y = T$ in the complex plane. We refer to this line as the *orbital line*, abbreviated as OL. Then the *Subset Sum Problem*, is to determine the number of points of the *point set* that lie on the OL.

DEFINITION 2.4. (*Extreme Points*) The points 0 and C_n which belong to the point set are referred to as the extreme points, since they correspond to the smallest and largest subset sums respectively. These two extreme points correspond to the binary vectors

$$\mathbf{0} := (0, 0, 0, \dots, 0, 0) \\ \text{and } \mathbf{1} := (1, 1, 1, \dots, 1, 1).$$

DEFINITION 2.5. For a given non-negative integer r (also referred to as index), where $r < 2^n$, we have an associated subset sum $\sigma_{\mathbf{a}}(r)$, defined by

$$\sigma_{\mathbf{a}}(r) := \sum_{u=1}^n k_u a_u, \quad \text{with } k_u \in \{0, 1\},$$

where $k_n k_{n-1} \dots k_2 k_1$ is the binary expansion of r . Note that when r is a small number that doesn't require n digits, we still use n digits, and the additional digits are treated as zeros.

2.1 A running example

To illustrate the main ideas of the paper, in particular the geometric structure, we use a running example which hereafter will be referred to as the *main example*.

EXAMPLE 2.6. (*Main Example*) This example is an instance of the Subset Sum Problem with $n = 9$ and $m = 18$,

$$\mathbf{a} = (43196, 84912, 106109, 119107, 119758, 121761, 125743, 131627, 147792)$$

and $T = 663708$. This instance has an unique solution with index equal to 365 whose binary form is $[365]_{\text{binary}} = 101101101$, and it can be verified that $663708 = 147792 + 125743 + 121761 + 119107 + 106109 + 43196$. While the algorithm for the Subset Sum Problem (to be discussed later) works for any positive (n, m) , we cannot illustrate the notions of the paper for large n without visual difficulty. We found that this is a good sized example to show the point set and paths through them.

2.2 Subspaces

The point set S_n is the complete solution space for SSP. However, it will be beneficial to look at the subspaces as well, which are defined below.

DEFINITION 2.7. (*Subspace*) For the given instance $\mathbf{a}$, a subspace of dimension $k \in [n]$ corresponds to the sequence $(a_1, a_2, \dots, a_k)$ with associated point set $2^{\mathbf{c}_k}$, where $\mathbf{c}_k = \{c_1, \dots, c_k\}$. We denote the point set of subspace k by S_k . Thus, we have n subspaces corresponding to the n possible values of k . All points in a subspace S_k lie in the interval box $[0, B_k] \times [0, A_k]$ in the Cartesian plane.

With the above definition, we can make unambiguous statements such as:

- (i) The point set $2^{\mathbf{c}}$ can be seen to be identical to two copies of subspace $2^{\mathbf{c}_{n-1}}$; one located at the origin, and the second located at c_n .
- (ii) Similarly, the point set $2^{\mathbf{c}}$ is the union of 2^{n-k} translates of the subspace $2^{\mathbf{c}_k}$.

2.3 Additional Definitions

We frequently refer to some progressions, defined below, that are utilized as some special cases of the *Subset Sum Problem*.

DEFINITION 2.8. (*Arithmetic Progression (AP)*) An instance $\mathbf{a}$ is said to be a n term arithmetic progression, denoted by $\mathbf{a} \in AP(n, k_1, k_2)$, or even simply by $\mathbf{a} \in AP$, if and only if there exist positive integer constants $k_1, k_2 \in \mathbb{N}$ such that

$$a_i = k_1 + (i - 1)k_2, \text{ for all } i \in [n].$$

DEFINITION 2.9. (*Constant Progression (CP)*) An instance $\mathbf{a}$ is said to be a n term constant progression, denoted by $\mathbf{a} \in CP(n, k_1)$, or even simply by $\mathbf{a} \in CP$, if and only if there is a positive integer constant $k_1 \in \mathbb{N}$ such that

$$a_i = k_1 \text{ for all } i \in [n].$$

DEFINITION 2.10. (*Geometric Progression (GP)*) An instance $\mathbf{a}$ is said to be a n term geometric progression, denoted by $\mathbf{a} \in GP(n, r)$, or even simply by $\mathbf{a} \in GP$, if and only if there exists a positive integer constant $r > 1$ such that

$$a_i = r^{i-1}, \text{ for all } i \in [n].$$

DEFINITION 2.11. (*Dissociated set*) Given an instance of SSP $(\mathbf{a}, T)$ we say that $\mathbf{a}$ is a dissociated set if and only if all the subset sums $2^{\mathbf{a}}$ are distinct.

3 Filling boxes, and paths from 0 to 1

Given an instance $\mathbf{a}$ of size n , the point set forms some configuration in the Cartesian plane, and this configuration changes with each instance. The configuration can take a wide range of shapes: for example, when $\mathbf{a} = \mathbf{b}$, all the points will be on a straight line; when $\mathbf{a} \in CP$, i.e. all elements are equal to a fixed positive integer, the points form $n + 1$ distinct levels along y ; and for a random instance they appear spread out in the whole box $[0, B_n] \times [0, A_n]$. This makes it harder to study and characterize the structure of the point set.

Furthermore, let $x_1, x_2 \in [B_n]$ be two random positive integers, and let $y_1 = \sigma_{\mathbf{a}}(x_1)$ and $y_2 = \sigma_{\mathbf{a}}(x_2)$ be the corresponding subset sums. Then the relative positions of $z_1 = (x_1, y_1)$ and $z_2 = (x_2, y_2)$ in the y coordinate can vary with instance. If $y_1 > y_2$ for one instance, it is possible to have $y_1 < y_2$ for another instance of the same size. Given these observations, consider the following question:

QUESTION 3.1. *For a fixed (n, m) , it can be seen that there are at most 2^{mn} possible instances of $\mathbf{a}$. Is there a universal structure that captures all these instances, and enable us to make (non-trivial) statements that are true for any instance?*

The answer to this question, as the reader will see, is a *resounding yes!*

While the relative ordering of the point set changes with instance, there exist points where a relative ordering is preserved for any instance. There are two extreme points, namely $(0, 0)$ and (B_n, A_n) where the second point is always at least as large in y coordinate as the first point for all instances of size n . Recall that these two extreme points correspond to the binary vectors $\mathbf{0}$ and $\mathbf{1}$. We characterize the structure of the solution space through the notion of non-decreasing paths between $\mathbf{0}$ and $\mathbf{1}$, which is described next.

DEFINITION 3.2. (*Path*) A path ψ of length N , from $\mathbf{0}$ to $\mathbf{1}$ is a sequence of binary vectors $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)$ where $\mathbf{x}_i \in \{0, 1\}^n$ for $i \in [N]$, such that $\mathbf{0} = \mathbf{x}_1$, $\mathbf{1} = \mathbf{x}_N$, and no vector is repeated. It is clear that $2 \leq N \leq 2^n$.

DEFINITION 3.3. (*Non-Decreasing Path (NDP)*) A path ψ from $\mathbf{0}$ to $\mathbf{1}$ is said to be a non-decreasing path, or simply a NDP, if and only if the subset sums along the path, are non-decreasing for **every** instance $\mathbf{a}$ of size n .

Consider the following simple example, to see that characterization of NDPs is not obvious even for simple cases.

EXAMPLE 3.4. Consider the simple case when $n = 3$. The sequence of indices for the power set in increasing order are $\mathfrak{R}(S_n) = (0, 1, 2, \dots, 7)$. If we write the corresponding sums, we get

$$\mathfrak{S}(S_n) = \{0, a_1, a_2, a_1 + a_2, a_3, a_3 + a_1, a_3 + a_2, a_3 + a_2 + a_1\}.$$

Does $\mathfrak{S}(S_n)$ form a non-decreasing sequence of sums for every $\mathbf{a}$? A moment's thought reveals that the answer is no, since a_3 need not be greater than or equal to $a_1 + a_2$. In

other words, the location of a_3 in the sequence is uncertain:

$$\begin{aligned} & \{0, a_1, a_2, a_1 + a_2, \boxed{a_3}, a_3 + a_1, a_3 + a_2, a_3 + a_2 + a_1\} \quad \text{if } a_3 > a_1 + a_2 \\ \text{or } & \{0, a_1, a_2, \boxed{a_3}, a_1 + a_2, a_3 + a_1, a_3 + a_2, a_3 + a_2 + a_1\} \quad \text{if } a_3 < a_1 + a_2. \end{aligned}$$

For the special case, when $a_3 = a_1 + a_2$, both the sequences will be NDPs.

Why do we care about NDPs? Since NDPs are independent of the instance, they provide a basis for the *universal geometric structure* that we seek. By definition, a NDP ψ remains non-decreasing for all instances of $\mathbf{a}$ that satisfy $a_1 \leq a_2 \leq \dots \leq a_n$. This suggests that the “non-decreasing property” should be independent of the magnitude of elements of $\mathbf{a}$. Alternatively, the non-decreasing property depends only on the binary vectors (indices) along the path. To be able to systematically characterize the NDPs, we consider a *model for filling n boxes with balls*, which is described next.

3.1 A model for filling n boxes

We are given n boxes, that abut each other with tops open and arranged in a line from left to right. We also have a bag of n balls to the immediate right of the rightmost box. Initially all the boxes are empty, and we associate this *empty state* with $\mathbf{0}$. We label the boxes as $n, n-1, \dots, 2, 1$ from left to right. The bag containing the balls is thought of as being at location 0. We have a machine that follows a single rule: *move a ball from location k to location $k+1$ provided it is empty, where $k \in [0, n-1]$* . For example, in the first step, the machine can only take a ball from the bag and put in box 1. In the second step, the machine picks up the ball in box 1, and displaces it to box 2. In the third step, the machine can either displace the ball in box 2 to box 3, or pick a new ball from the bag and put in box 1. At each step, there is a unit displacement from one box to the box on left (provided it is empty), or from the bag to box 1. Note that the displacements are always to the left by unit distance, and never to the right. Also we cannot displace a ball to a position that is not empty. This process is repeated until all the boxes are filled. The state where all the n boxes are filled is associated with $\mathbf{1}$.

The following proposition is straightforward.

PROPOSITION 3.5. *To go from the empty state to the filled state, we need exactly $n(n+1)/2$ unit displacements.*

DEFINITION 3.6. (*state and sequence of states*) We refer to the configuration of balls in the n boxes at any time step as a state. By associating a 0 with an empty box, and a 1 with a filled box, a state corresponds to a binary vector of length n . Let $\pi := (\pi_0, \pi_1, \pi_2, \dots, \pi_{n(n+1)/2})$ be the sequence of states in filling the boxes, where $\pi_k \in \{0, 1\}^n$ for $k \in \{0, 1, \dots, n(n+1)/2\}$, π_0 is the empty state, and $\pi_{n(n+1)/2}$ is the filled state.

PROPOSITION 3.7. *Any sequence of states obeying the rules for filling the boxes corresponds to a NDP.*

Proof. Since we always move a ball from a lower indexed box to a higher indexed box, it amounts to replacing a_k by a_{k+1} for $k \in [n-1]$, or adding a a_1 to the subset sum. Since $\mathbf{a}$ satisfies $a_1 \leq a_2 \leq \dots \leq a_n$, the subset sums corresponding to the states are non-decreasing, and the conclusion follows. $\square$

3.2 Two complementary ways of filling boxes

It is clear that, as long as the balls keep moving to the left, one displacement at a time, any sequence of choices can be used to fill all the boxes, and the sequence of corresponding binary vectors will result in non-decreasing subset sums. Instead of randomly choosing which ball to move at a time, we will focus on two particular ways, temporarily labeled LB and HB:

- (LB) At each time step, choose the ball at the lowest *possible* bit position that *can* be moved one unit to the left.
- (HB) At each time step, choose the ball at the highest *possible* bit position that *can* be moved one unit to the left.

An example will clarify these two ways of filling all boxes.

EXAMPLE 3.8. Let $n = 4$. The sequence obtained using LB is on the left, and the sequence obtained using HB is in the center. We use 0 to represent an empty box, and a 1 for a filled box. The balls enter from the right, starting in the rightmost box. Notice that both sequences have a length of $1 + 4 \times 5/2 = 11$. The third table, labeled RC, uses a random choice.

		LB				HB					RC	
0	0	0	0		0	0	0	0		0	0	0
0	0	0	1		0	0	0	1		0	0	0
0	0	1	0		0	0	1	0		0	0	1
0	0	1	1		0	1	0	0		0	1	0
0	1	0	1		1	0	0	0		0	1	0
0	1	1	0		1	0	0	1		1	0	0
0	1	1	1		1	0	1	0		1	0	1
1	0	1	1		1	1	0	0		1	0	1
1	1	0	1		1	1	0	1		1	1	0
1	1	1	0		1	1	1	0		1	1	1
1	1	1	1		1	1	1	1		1	1	1

We are interested in LB and HB, as they turn out to be important for the SSP. For an arbitrary n , we can generate the LB/HB sequence algorithmically, simply by applying the rule for $n(n+1)/2$ steps, starting with the zero vector.

LEMMA 3.9. Let rule LB be used to fill k boxes. If we add another box on the left, add a new ball to the bag, and continue using rule LB, then we get $k+1$ new elements into the sequence.

Proof. For the problem with size k , the last element in the sequence has the following binary form of length k .

$$\begin{array}{ccccccc} x_k & x_{k-1} & x_{k-2} & \dots & x_3 & x_2 & x_1 \\ 1 & 1 & 1 & \dots & 1 & 1 & 1 \end{array}$$

Now the next elements will be, using the rules defined above,

$$\begin{array}{ccccccc} x_{k+1} & x_k & x_{k-1} & x_{k-2} & \dots & x_3 & x_2 & x_1 \\ \boxed{0} & \boxed{1} & \boxed{1} & \boxed{1} & \dots & \boxed{1} & \boxed{1} & \boxed{1} \\ 1 & 0 & 1 & 1 & \dots & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & \dots & 1 & 1 & 1 \\ \vdots & & & & & & & \\ 1 & 1 & 1 & 1 & \dots & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & \dots & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & \dots & 1 & 1 & 1 \end{array}$$

Since there are k moves of 1s and an added 1 (from bag) for the last element, we get $k + 1$ new elements going from problem size k to $k + 1$. $\square$

3.3 Formulae to generate NDS

It turns out that the non-decreasing sequences described algorithmically above can be generated by formulae.

DEFINITION 3.10. For a non-negative integer k , let

$$\vartheta(k) := \lfloor \frac{-1 + \sqrt{1 + 8k}}{2} \rfloor$$

where $\lfloor x \rfloor$ denotes the largest integer less than or equal to x . Thus, $\vartheta(0) = 0, \vartheta(1) = 1, \vartheta(2) = 1, \vartheta(3) = 2$ and so on.

THEOREM 3.11. The non-decreasing sequence corresponding to using rule LB, denoted by ϕ_n , can be generated by the formula²

$$\phi_n(k) := 2^{1+\vartheta(k)} - 2^{\vartheta(k)-k+\frac{\vartheta(k)(1+\vartheta(k))}{2}} - 1, \quad (3.1)$$

where $k \in [0, \frac{n(n+1)}{2}]$.

Proof. The proof is given by induction on n .

(Base Step) Consider first the case when $n = 2$. We have a total of $\frac{2 \cdot 3}{2} = 3$ elements in the sequence. Using the formula for ϑ defined above and letting $i = 0, 1, 2, 3$ we have

$$\vartheta(0) = 0, \vartheta(1) = 1, \vartheta(2) = 1, \text{ and } \vartheta(3) = 2.$$

²While this formula was discovered independently, it is not new. In a different context, the sequence ϕ_n is known [see OEIS at <https://oeis.org/A089633>].

Now the ϕ values are

$$\begin{aligned}\phi_2(0) &= 2^{(1+0)} - 2^{(0-0+\frac{0+1}{2})} - 1 = 0 \equiv [0, 0]_{2,2} \\ \phi_2(1) &= 2^{(1+1)} - 2^{(1-1+\frac{1+2}{2})} - 1 = 1 \equiv [0, 1]_{2,2} \\ \phi_2(2) &= 2^{(1+1)} - 2^{(1-2+\frac{1+2}{2})} - 1 = 2 \equiv [1, 0]_{2,2} \\ \phi_2(3) &= 2^{(1+2)} - 2^{(2-3+\frac{2+3}{2})} - 1 = 3 \equiv [1, 1]_{2,2}\end{aligned}$$

where the corresponding binary representation is shown at the right. (On the right, the first subscript means base 2, and the second one means we are using 2 digits.) So the assertion is true for $n = 2$.

(Inductive Step) Now, assume that it is true for $n = k > 2$. For this sequence, the last element has an index equal to $\frac{k(k+1)}{2}$. When we increment n to be equal to $k + 1$, we are interested in the range

$$\frac{k(k+1)}{2} + 1, \frac{k(k+1)}{2} + 2, \dots, \frac{(k+1)(k+2)}{2}$$

The corresponding ϑ values are

$$\begin{aligned}\vartheta\left(\frac{k(k+1)}{2} + j\right) &= \left\lfloor \frac{-1 + \sqrt{(2k+1)^2 + 8j}}{2} \right\rfloor \\ &= \begin{cases} k & \text{if } j = 1, 2, \dots, k \\ k+1 & \text{if } j = k+1 \end{cases}\end{aligned}$$

Therefore,

$$\phi_{k+1}\left(\frac{k(k+1)}{2} + j\right) = \begin{cases} 2^{1+k} - 2^{k-(\frac{k(k+1)}{2}+j)+\frac{k(k+1)}{2}} - 1 & \text{if } j \in \{1, 2, \dots, k\} \\ 2^{2+k} - 2^{1+k} - 1 & \text{if } j = k+1 \end{cases}$$

First consider the case $j \in [k]$: it can be seen that $2^{1+k} - 2^{k-(\frac{k(k+1)}{2}+j)+\frac{k(k+1)}{2}} - 1 = 2^{1+k} - 2^{k-j} - 1$ after canceling some terms. For $j = 0$, we get $2^{1+k} - 2^k - 1 = 2^k - 1$ whose binary form consists of k 1s. This is the final element for $n = k$. When $j = 1$, we get $2^{1+k} - 2^{k-1} - 1 = 2^k + 2^{k-1} - 1$ as the 1 in the k th place moves to $k+1$ th place leaving a zero at the k th place. As j is increased further, the zero location propagates all the way to the 1st position. This position gets filled when $j = k+1$. Thus, the bit transitions happen as in Lemma 3.9 and we have a non-decreasing sequence. Hence it is true for all n . $\square$

COROLLARY 3.12. *The sequence corresponding to using rule HB, denoted by φ_n , can be described by the formula*

$$\varphi_n(k) = \phi_n^c(n(n+1)/2 - k), \text{ for } k = 0, 1, \dots, n(n+1)/2, \quad (3.2)$$

where $\phi_n^c(k)$ is the binary 1s complement of $\phi_n(k)$, i.e.,

$$\phi_n^c(k) = 2^n - 1 - \phi_n(k) = 2^n - 2^{(1+\vartheta(k))} + 2^{\vartheta(k)-k+\frac{\vartheta(k)(\vartheta(k)+1)}{2}}. \quad (3.3)$$

Proof. Since ϕ_n is non-decreasing, the reverse of this sequence denoted $REV(\phi_n)$ is non-increasing. Now consider the 1s complement of $REV(\phi_n)$, denoted by $REV(\phi_n^c)$. Clearly, this is non-decreasing. Since the sequence ϕ_n corresponds to rule LB, where we move the ball in lowest possible bit position, when we apply the reversal and the 1s complement, it becomes the highest possible bit position. As a result, it corresponds to filling boxes using rule HB. $\square$

As a consequence of Theorem 3.11 and Corollary 3.12, we have the following result, due to their complementary nature.

PROPOSITION 3.13. *For any $j \in [n]$, and $k \in [0, j(j+1)/2]$, we have the identities*

- (i) $\phi_j(k) + \varphi_j(\frac{j(j+1)}{2} - k) = B_j$, and
- (ii) $\sigma_{\mathbf{a}}(\phi_j(k)) + \sigma_{\mathbf{a}}(\varphi_j(\frac{j(j+1)}{2} - k)) = A_j$.

The two families of complementary sequences are shown below for clarity. The first family of sequences labeled $\phi_k, k = 1, 2, \dots, n$ for varying problem sizes is

$$\begin{aligned}
\phi_1 &= \{0, 1\} \\
\phi_2 &= \{0, 1, 2, 3\} \\
\phi_3 &= \{0, 1, 2, \boxed{3}, 5, 6, 7\} \\
\phi_4 &= \{0, 1, 2, \boxed{3, 5, 6, 7, 11}, 13, 14, 15\} \\
\phi_5 &= \{0, 1, 2, \boxed{3, 5, 6, 7, 11, 13, 14, 15, 23, 27}, 29, 30, 31\} \\
\phi_6 &= \{0, 1, 2, \boxed{3, 5, 6, 7, 11, 13, 14, 15, 23, 27, 29, 30, 31, 47, 55, 59}, 61, 62, 63\} \\
&\dots \\
\phi_n &= \{0, 1, 2, \boxed{3, 5, \dots, 2^n - 5}, 2^n - 3, 2^n - 2, 2^n - 1\}
\end{aligned}$$

The second family of sequences labeled $\varphi_k, k = 1, 2, \dots, n$ is

$$\begin{aligned}
\varphi_1 &= \{0, 1\} \\
\varphi_2 &= \{0, 1, 2, 3\} \\
\varphi_3 &= \{0, 1, 2, \boxed{4}, 5, 6, 7\} \\
\varphi_4 &= \{0, 1, 2, \boxed{4, 8, 9, 10, 12}, 13, 14, 15\} \\
\varphi_5 &= \{0, 1, 2, \boxed{4, 8, 16, 17, 18, 20, 24, 25, 26, 28}, 29, 30, 31\} \\
\varphi_6 &= \{0, 1, 2, \boxed{4, 8, 16, 32, 33, 34, 36, 40, 48, 49, 50, 52, 56, 57, 58, 60}, 61, 62, 63\} \\
&\dots \\
\varphi_n &= \{0, 1, 2, \boxed{4, 8, \dots, 2^n - 4}, 2^n - 3, 2^n - 2, 2^n - 1\}
\end{aligned}$$

Note that while $\phi_1 = \varphi_1$ and $\phi_2 = \varphi_2$, other pairs are different in some elements. The differences are *boxed* to make them clear. The sequences $\phi_n(i)$ and $\varphi_n(i)$ for all

$n > 2$ have exactly six points in common. These six points in decimal values are $0, 1, 2, 2^n - 3, 2^n - 2$, and $2^n - 1$. As a result, the pair $\{\phi_n, \varphi_n\}$ of sequences together contain $2(1 + \frac{n(n+1)}{2}) - 6 = n(n+1) - 4$ unique points for all $n > 2$.

3.4 From sequences to curves

DEFINITION 3.14. For a given positive integer $j \in [n]$, let p_j be the sequence of complex numbers defined by

$$p_j(k) = \phi_j(k) + \iota \sigma_{\mathbf{a}}(\phi_j(k)) \quad \text{for } k = 0, 1, 2, \dots, \frac{j(j+1)}{2}. \quad (3.4)$$

Similarly, let q_j be the sequence of complex numbers defined by

$$q_j(k) = \varphi_j(k) + \iota \sigma_{\mathbf{a}}(\varphi_j(k)) \quad \text{for } k = 0, 1, 2, \dots, \frac{j(j+1)}{2}. \quad (3.5)$$

The real part corresponds to the index, and the imaginary part corresponds to the subset sum associated with the index, for both sequences.

By joining adjacent points of p_j by straight lines, we obtain a continuous piecewise-linear curve $p_j(t)$ in the complex plane, with a natural parametrization, $t \in [0, \frac{j(j+1)}{2}]$. Similarly for $q_j(t)$.

DEFINITION 3.15. (Vertex Points) The integral points of p_j and q_j , defined by Equation (3.4) and Equation (3.5) respectively, are called vertex points to distinguish them from other points on the piecewise-linear curves. In other words, a point $(x, y) \in p_j$ (or q_j) is said to be a vertex point if and only if it belongs to the point set S_n . Clearly, p_j has exactly $1 + \frac{j(j+1)}{2}$ vertex points, and so does q_j .

3.5 Reflection Symmetry

By the complementarity of the indices (ϕ_k and φ_k in bit form) discussed above, it follows that for any $t \in [0, N]$ where $N = \frac{k(k+1)}{2}$, we have

$$p_k(t) + q_k(N - t) = C_k. \quad (3.6)$$

This can be written as

$$p_k(t) - \frac{C_k}{2} = \frac{C_k}{2} - q_k(N - t)$$

which has the form of an equation representing *Point Reflection Symmetry* of these two curves about the point $\frac{C_k}{2}$. Thus the complementary curves $p_k(t)$ and $q_k(t)$ satisfy a *Point Reflection Symmetry*. An example is shown in Figure 1 below.




Figure 1: Point Reflection Symmetry for complementary curves p_6 (blue) and q_6 (cyan) of the main example. The symmetric pairs of points are connected by a dotted line.

3.6 The curves p_k and q_k and coverage

Here, we provide the geometrical intuition, by showing the pair of curves p_k, q_k and the subspace S_k as k ranges in $\{1, \dots, 9\}$ using the main example. These are shown in Figure 2, where the p_k curve is in blue, and the q_k curve is in cyan. Some observations follow:

- [Cases $n = 1$ and $n = 2$] For the case $n = 1$ the solution space is $\{0, 1\}$, and the sequences ϕ_1 and φ_1 are both identical and equal to $\{0, 1\}$. For $n = 2$, we have 4 subset sums corresponding to indices $\{0, 1, 2, 3\}$. The sequences, $\phi_2 = \{0, 1, 2, 3\} = \varphi_2$ are identical. In these trivial cases, just one sequence of indices covers the solution space.
- [Case $n = 3$] For $n = 3$, the indices of the solution space are $\{0, 1, 2, 3, 4, 5, 6, 7\}$. The sequences $\phi_3 = \{0, 1, 2, 3, 5, 6, 7\}$ and $\varphi_3 = \{0, 1, 2, 4, 5, 6, 7\}$ are distinct. Note that the index 3 is in ϕ_3 but not in φ_3 . Similarly, the index 4 is in φ_3 but not in ϕ_3 . However, together they cover the solution space.
- [Case $n = 4$] For $n = 4$, we have 16 elements in the solution space. It is easy to verify that the 16 indices are covered by $\phi_4 = \{0, 1, 2, 3, 5, 6, 7, 11, 13, 14, 15\}$ and $\varphi_4 = \{0, 1, 2, 4, 8, 9, 10, 12, 13, 14, 15\}$ together, i.e., $\phi_4 \cup \varphi_4 = S_4$.
- [Case $n = 5$] For $n = 5$, the indices are $\{0, 1, 2, \dots, 31\}$. The sequences ϕ_5 and φ_5 are $\{0, 1, 2, 3, 5, 6, 7, 11, 13, 14, 15, 23, 27, 29, 30, 31\}$ and $\{0, 1, 2, 4, 8, 16, 17, 18, 20, 24, 25, 26, 28, 29, 30, 31\}$ respectively. However, notice that these two sequences,

even together do not cover the solution space. In particular, the indices $\{9, 10, 12, 19, 21, 22\}$ are missing.

- [Cases $n > 5$] For $n > 5$, the number of elements of the subspace that are not covered by p_n and q_n will increase with n . We will show later on that by including additional curves (formed by certain operations on the p_n and q_n curves) we can get complete coverage of the solution space. This will be discussed in detail in the next section, after developing the suitable framework.

4 Links, Chains and Transforms

In the previous section we used the Filling Boxes Model to characterize two particular NDPs namely $p_n(k)$ and $q_n(k)$ where $k \in [0, \frac{n(n+1)}{2}]$. Here we show an alternate characterization of these NDPs using the differences of adjacent elements of the sequence **c**. Furthermore, both characterizations are related in a way analogous to *differentiation* and *integration*.

DEFINITION 4.1. (*Set of Differences D and links*) Let $d_j := c_j - c_{j-1}$, for $j \in [n]$ where $c_0 := 0$, and let $D := \{d_1, d_2, \dots, d_n\}$ be the set of complex numbers corresponding to the differences. The complex number d_j can also be viewed as a directed line segment from the origin to the point $(b_j - b_{j-1}) + \iota(a_j - a_{j-1})$ for each $j \in [n]$. When referring to the line segment, we call it as a link but still denote it as d_j for simplicity.

Since the maximum size of each element in the given instance is bounded by 2^m , it follows that $\Im(d_j) < 2^m$ for all $j \in [n]$.

For any positive integer $j \in [n]$, we can write c_j as a telescoping sum:

$$c_j = d_1 + \dots + d_j.$$

PROPOSITION 4.2. Let $j \in [n]$ be a positive integer. Then for any non-negative integer $k \in [0, \frac{j(j+1)}{2})$ we have

$$\begin{aligned} p_j(k+1) - p_j(k) &= d_r, \\ \text{and } q_j(k+1) - q_j(k) &= d_s \end{aligned}$$

for some $r, s \in [n]$.

Proof. Recall that for any non-negative integer k , we have $p_j(k) = \phi_j(k) + \iota\sigma_{\mathbf{a}}(\phi_j(k))$. In the sequence ϕ_j , as we go from one element to the next element, the Filling-Boxes model shows that we are displacing a ball from a lower bit position to an immediate higher bit position. Thus, the difference between adjacent indices has the form $b_r - b_{r-1}$ for some suitable $r \in [n]$. The corresponding change in the subset sum, is clearly $a_r - a_{r-1}$. Thus, from Definition 3.14 it follows that $p_j(k+1) - p_j(k) = d_r$. Similarly for q_j whose indices are defined by the sequence φ_j . $\square$




Figure 2: Illustration of coverage of subspaces $S_1, S_2, \dots, S_9$, by complementary curves p_k and q_k . In each case the p_k curve is in blue and the q_k curve in cyan. In S_1 and S_2 the complementary curves are identical and overlap. In S_3 , the curves p_3 and q_3 are distinct in one point, and together they cover the space. In S_4 , the curves p_4 and q_4 are distinct in more points, and together they cover the space. In S_5 , the curves p_5 and q_5 are distinct in several points but do not cover the space. There are several isolated points. This pattern continues for higher subspaces.

DEFINITION 4.3. (*Chain*) A chain is simply an ordered concatenation of two or more elements of D , where the operation of concatenation is denoted by $\oplus$ which separates the given numbers. A chain is indicated by the symbol ∂ preceding an alphabetical character. For example, the chain ∂L which is the concatenation of d_i and d_j , for some $i, j \in [n]$, is written as $\partial L = d_i \oplus d_j$. We refer to the k th element of a chain by $\partial L(k)$. Similarly for $i, j \in [n]$ with $i < j$ the expression $\partial L[i, j]$ indicates the sequence of elements (subchain) starting at index i and ending at index j .

With these definitions, we can associate the curves p_n and q_n to chains as shown in the proposition below. The elements in the chain are *boxed* to disclose the pattern clearly.

PROPOSITION 4.4. The sequence of differences between adjacent points in the sequence p_n , corresponds to the chain ∂p_n , where

$$\partial p_n = \boxed{d_1} \oplus \boxed{d_2 \oplus d_1} \oplus \boxed{d_3 \oplus d_2 \oplus d_1} \oplus \dots \oplus \boxed{d_n \oplus d_{n-1} \oplus \dots \oplus d_1}. \quad (4.1)$$

The sequence of differences between adjacent points in the sequence q_n corresponds to the chain ∂q_n , where

$$\partial q_n = \boxed{d_1 \oplus d_2 \oplus \dots \oplus d_n} \oplus \boxed{d_1 \oplus \dots \oplus d_{n-1}} \oplus \dots \oplus \boxed{d_1 \oplus d_2} \oplus \boxed{d_1}. \quad (4.2)$$

Proof. We give the proof by induction on n .

(Base Step) For $n = 1$, the sequence of indices $\phi_1 = (0, 1)$ and sums $\sigma_{\mathbf{a}}(\phi_1) = (0, a_1)$. The chain has only one element, i.e., $(b_1 + \iota a_1) - (b_0 + \iota a_0) = 1 + \iota a_1$ which is defined by d_1 . Thus $\partial p_1 = (d_1)$.

(Inductive Step) By hypothesis suppose it is true for $n = k$. So we have,

$$\partial p_k = \boxed{d_1} \oplus \boxed{d_2 \oplus d_1} \oplus \boxed{d_3 \oplus d_2 \oplus d_1} \oplus \dots \oplus \boxed{d_k \oplus d_{k-1} \oplus \dots \oplus d_1}.$$

For $n = k + 1$, let's refer to the Filling-Boxes model. We have an empty box on the left and an extra ball in the bag. Also, all the k boxes to the left of the bag are filled. The next sequence of displacements happen as indicated in Lemma 3.9, which correspond to

$$(d_{k+1}, d_k, \dots, d_2, d_1).$$

Thus we have

$$\begin{aligned} \partial p_{k+1} &= \partial p_k \oplus \boxed{d_{k+1} \oplus d_k \oplus \dots \oplus d_1} \\ &= \boxed{d_1} \oplus \boxed{d_2 \oplus d_1} \oplus \boxed{d_3 \oplus d_2 \oplus d_1} \oplus \dots \oplus \boxed{d_{k+1} \oplus d_k \oplus \dots \oplus d_1}. \end{aligned}$$

and the claim follows. By recalling that φ_n is the bit complement of ϕ_n , it follows that the chain corresponding to q_n is the reverse of that of p_n . $\square$

COROLLARY 4.5. Let $N = \frac{n(n+1)}{2}$. Then for any $k \in [N]$ we have

$$\partial p_n(k) = \partial q_n(N + 1 - k).$$

PROPOSITION 4.6. To each real number $t \in [0, n(n+1)/2]$, the unique point $p_n(t)$ can be determined as

$$p_n(t) = \sum_{j=1}^{\lfloor t \rfloor} \partial p_n(j) + (t - \lfloor t \rfloor) \partial p_n(\lfloor t \rfloor + 1)$$

where, in the first term, the empty sum is treated as zero. Similarly,

$$q_n(t) = \sum_{j=1}^{\lfloor t \rfloor} \partial q_n(j) + (t - \lfloor t \rfloor) \partial q_n(\lfloor t \rfloor + 1).$$

Proof. First suppose that $t \in \mathbb{N}$. Then $p_n(t)$ is simply the sum of the first t elements of the chain ∂p_n which is nothing but a telescoping sum. If t is not an integer, then we get the extra term on the right corresponding to a partial traversal on a link, since the curve $p_n(t)$ is a piecewise-linear curve. $\square$

With some abuse of notation, the above proposition can be written more elegantly as follows.

PROPOSITION 4.7. For any positive integer $j \in [n]$ and for any real $t \in [0, \frac{j(j+1)}{2}]$

$$p_j(t) = \int_0^t \partial p_j(t') dt'$$

$$\text{and } q_j(t) = \int_0^t \partial q_j(t') dt'.$$

Thus we have two equivalent ways of representing the curve: the curve itself (i.e. a sequence of points), and a chain (i.e. a sequence of line segments) where to get the coordinates of a point we take a path integral. Thus, informally, a chain is like the derivative of the curve with respect to parameter t , while the curve is like the integral of chain.

4.1 Elemental representation of chains

A clear pattern is evident in Equation (4.1) and Equation (4.2), as shown by *boxing* the elements, which leads to the definition of an *elemental chain*. This will simplify the notation and make the description easy as we move forward.

DEFINITION 4.8. (*Elemental Chains*) We view the elements of D as building blocks of elements of $\mathbf{c}$. In particular, to each $c_i \in \mathbf{c}$, we associate a chain ∂c_i defined by

$$\partial c_i = d_1 \oplus d_2 \oplus \dots \oplus d_i. \quad (4.3)$$

We also define the reverse chain (with a hat symbol)

$$\partial \hat{c}_i = d_i \oplus d_{i-1} \oplus \dots \oplus d_1. \quad (4.4)$$

We refer to ∂c_i and $\partial \hat{c}_i$ as elemental chains which are reverse of each other.

With this definition, the chain ∂p_n can be written compactly as

$$\partial p_n = \partial \hat{c}_1 \oplus \partial \hat{c}_2 \oplus \dots \oplus \partial \hat{c}_n. \quad (4.5)$$

Similarly, the chain ∂q_n can be written as,

$$\partial q_n = \partial c_n \oplus \partial c_{n-1} \dots \oplus \partial c_1. \quad (4.6)$$

OBSERVATION 4.9. *The following observations can be made with regard to the elemental representation of chains ∂p_n and ∂q_n .*

- (i) *In ∂p_n , if we consider only the first elemental chain, we see that it corresponds to ∂p_1 . In general, if we consider the first k elemental chains, for any $k \in [n]$, the chain corresponds to ∂p_k . Thus, the chain ∂p_n contains all the sub-chains $\partial p_1, \partial p_2, \dots, \partial p_{n-1}$.*
- (ii) *In ∂q_n , if we consider only the last elemental chain, we see that it corresponds to ∂q_1 . In general, if we consider the last k elemental chains, for any $k \in [n]$, the chain corresponds to ∂q_k . Thus, the chain ∂q_n contains all the sub-chains $\partial q_1, \partial q_2, \dots, \partial q_{n-1}$.*

The above observations can be formally stated as:

PROPOSITION 4.10. *Let $k \in [n]$ be a positive integer. Then the chain ∂p_k contains $\partial p_1, \dots, \partial p_{k-1}$ as subchains. Similarly, the chain ∂q_k contains $\partial q_1, \dots, \partial q_{k-1}$ as subchains.*

DEFINITION 4.11. *(Partial Sums of a Chain) For a positive integer N , let $\partial L = l_1 \oplus \dots \oplus l_N$ be a chain of length N . We denote the sum of the first i elements of ∂L by $\sigma(\partial L, i)$ and the sum of all elements of ∂L simply by $\sigma(\partial L)$. We also define the empty sum to be 0, i.e., $\sigma(\partial L, 0) = 0$. Thus, $\sigma(\partial L, 2) = l_1 + l_2$ and $\sigma(\partial L) = l_1 + \dots + l_N$.*

DEFINITION 4.12. *(Valid Subset Sum) A complex number s is said to be a valid subset sum if and only if $s \in S_n$, i.e., s is an element of the power set.*

It is straight forward to see that, for any $k \in [n]$, all the partial sums of both ∂p_k and ∂q_k are indeed valid subset sums.

4.2 Transforms on chains

As noted above, the pair of complementary sequences p_n and q_n cover the solution space only for $n \leq 4$. For $n \geq 5$ we will have many missing points, and in order to cover them we need *new* sequences/chains that are distinct from ∂p_n and ∂q_n . In this section, we define certain transforms that enable us to form new chains from a given chain.

DEFINITION 4.13. *(Active Region) Let $\partial \psi$ be a chain consisting of n elemental chains. Then the active region of $\partial \psi$, denoted by $\alpha(\partial \psi)$, is defined as the longest interval of indices $[i, j] \subseteq [1, n]$ satisfying either $\partial \psi[i, j] = \partial p_k$ or $\partial \psi[i, j] = \partial q_k$, where $k = j + 1 - i$. The non-active part of a chain is said to be fixed.*

Thus, it readily follows that $\alpha(\partial p_n) = \alpha(\partial q_n) = [1, n]$. Below we define some transformations on chains which operate *only* on the active region of a given chain.

DEFINITION 4.14. (*Character set Λ*) Given a chain $\partial\psi$ with active region specified by $[i, j]$, the character set of $\partial\psi$ denoted by $\Lambda(\partial\psi)$, is defined as

$$\Lambda(\partial\psi) = [j + 1 - i].$$

Thus, the character set Λ of a chain $\partial\psi$ tells us that there exist subchains of either type ∂p_r or type ∂q_r within $\partial\psi$ for all $r = 1, 2, \dots, j + 1 - i$.

EXAMPLE 4.15. Let $\partial\psi$ be a chain consisting of seven blocks, described by

$$\partial\psi = \partial c_5 \oplus \partial c_3 \oplus \partial c_2 \oplus \partial c_1 \oplus \partial \hat{c}_4 \oplus \partial \hat{c}_6 \oplus \partial \hat{c}_7$$

Then, from the above definitions it follows that $\alpha(\partial\psi) = [2, 4]$ since these indices correspond to the longest subchain $\partial q_3 = \partial c_3 \oplus \partial c_2 \oplus \partial c_1$. Thus $\Lambda(\partial\psi) = [3]$ and we have the complete subchains $\partial q_1 = \partial c_1$, $\partial q_2 = \partial c_2 \oplus \partial c_1$ and $\partial q_3 = \partial c_3 \oplus \partial c_2 \oplus \partial c_1$.

On the other hand if we consider the chain ∂q_{10} we have $\alpha(\partial q_{10}) = [1, 10]$, and $\Lambda(\partial q_{10}) = [10]$.

DEFINITION 4.16. (*Local Character Transform λ_k*) Let $\partial\psi$ be a chain composed of n elemental chains, and active region specified by the interval $[i, j]$. Then given a positive integer $k \in [j + 1 - i]$, the local character transform denoted by $\lambda_k(\partial\psi)$ changes the active region as follows:

$$\alpha(\lambda_k(\partial\psi)) = \begin{cases} [i, i + k - 1] & \text{if } \partial\psi[i, j] \text{ is of the form } \partial p_{j+1-i} \\ [j + 1 - k, j] & \text{if } \partial\psi[i, j] \text{ is of the form } \partial q_{j+1-i}. \end{cases}$$

Informally, the local character transform picks up one of the available characters from the character set. Note that the local character transform doesn't change the chain; it simply changes the active region - the region that can undergo further transforms. Furthermore, the λ_k operation affects the leftmost blocks of p -type active region, but the rightmost blocks of q -type active region.

EXAMPLE 4.17. Consider the chain $\partial q_7 = \partial c_7 \oplus \partial c_6 \oplus \partial c_5 \oplus \partial c_4 \oplus \partial c_3 \oplus \partial c_2 \oplus \partial c_1$. Clearly, the active region is $[1, 7]$. If we apply λ_4 transform on this chain, we get

$$\lambda_4(\partial q_7) = \partial c_7 \oplus \partial c_6 \oplus \partial c_5 \oplus \boxed{\partial c_4 \oplus \partial c_3 \oplus \partial c_2 \oplus \partial c_1}$$

where the new active region is $[4, 7]$ whose elemental chains are boxed for clarity.

Similarly, consider the chain $\partial p_7 = \partial \hat{c}_1 \oplus \partial \hat{c}_2 \oplus \partial \hat{c}_3 \oplus \partial \hat{c}_4 \oplus \partial \hat{c}_5 \oplus \partial \hat{c}_6 \oplus \partial \hat{c}_7$. Clearly, the active region is $[1, 7]$. If we apply λ_5 transform on this chain, we get

$$\lambda_5(\partial p_7) = \boxed{\partial \hat{c}_1 \oplus \partial \hat{c}_2 \oplus \partial \hat{c}_3 \oplus \partial \hat{c}_4 \oplus \partial \hat{c}_5} \oplus \partial \hat{c}_6 \oplus \partial \hat{c}_7$$

where the new active region is $[1, 5]$ which is boxed for clarity.

DEFINITION 4.18. (Reversal Transform ρ_k) Let $\partial\psi$ be a chain with n blocks and active region specified by the interval $[i, j]$. Then given a positive integer $k = j + 1 - i$, the reversal transform denoted by $\rho_k(\partial\psi)$ reverses the order of elemental chains in the active region and reverses the order of elements in each elemental chain. Thus

$$\rho_k(\partial\psi[i, j]) = \begin{cases} \partial\hat{c}_1 \oplus \partial\hat{c}_2 \dots \oplus \partial\hat{c}_k & \text{if } \partial\psi[i, j] = \partial c_k \oplus \partial c_{k-1}, \dots \oplus \partial c_1 \\ \partial c_k \oplus \partial c_{k-1} \dots \oplus \partial c_1 & \text{if } \partial\psi[i, j] = \partial\hat{c}_1 \oplus \partial\hat{c}_2 \dots \oplus \partial\hat{c}_k. \end{cases}$$

In words, if the active region of the chain $\partial\psi$ originally corresponded to ∂p_k then after the ρ_k transform this portion of the chain becomes ∂q_k , and vice versa.

EXAMPLE 4.19. (Continuation of Example 4.17) Suppose we apply a ρ_4 transform on $\lambda_4(\partial p_7)$, we get:

$$\begin{aligned} \rho_4(\lambda_4(\partial q_7)) &= \rho_4(\partial c_7 \oplus \partial c_6 \oplus \partial c_5 \oplus \boxed{\partial c_4 \oplus \partial c_3 \oplus \partial c_2 \oplus \partial c_1}) \\ &= \partial c_7 \oplus \partial c_6 \oplus \partial c_5 \oplus \boxed{\partial\hat{c}_1 \oplus \partial\hat{c}_2 \oplus \partial\hat{c}_3 \oplus \partial\hat{c}_4}. \end{aligned}$$

Similarly if we apply a ρ_5 transform on $\lambda_5(\partial q_7)$ we get

$$\begin{aligned} \rho_5(\lambda_5(\partial p_7)) &= \rho_5(\boxed{\partial\hat{c}_1 \oplus \partial\hat{c}_2 \oplus \partial\hat{c}_3 \oplus \partial\hat{c}_4 \oplus \partial\hat{c}_5} \oplus \partial\hat{c}_6 \oplus \partial\hat{c}_7) \\ &= \boxed{\partial c_5 \oplus \partial c_4 \oplus \partial c_3 \oplus \partial c_2 \oplus \partial c_1} \oplus \partial\hat{c}_6 \oplus \partial\hat{c}_7. \end{aligned}$$

We can define a composite τ_k transform equivalent to a local character transform followed by a reversal transform as follows.

DEFINITION 4.20. (Reflection Transform τ_k) Let $\partial\psi$ be a chain with n elemental chains and active region specified by the interval $[i, j]$. Then given a positive integer $k \in [j + 1 - i]$, the reflection transform denoted by $\tau_k(\partial\psi)$ is defined as

$$\tau_k(\partial\psi) := \rho_k(\lambda_k(\partial\psi)).$$

Note that the τ_k transform operates on the first k elemental chains if the active region of $\partial\psi$ is of the type ∂p_{j+1-i} , and on the last k elemental chains if the active region is of the type ∂q_{j+1-i} .

EXAMPLE 4.21. With the composite reflection transform, the previous example can be written more simply as:

$$\begin{aligned} \tau_4(\partial q_7) &= \tau_4(\partial c_7 \oplus \partial c_6 \oplus \partial c_5 \oplus \partial c_4 \oplus \partial c_3 \oplus \partial c_2 \oplus \partial c_1) \\ &= \partial c_7 \oplus \partial c_6 \oplus \partial c_5 \oplus \boxed{\partial\hat{c}_1 \oplus \partial\hat{c}_2 \oplus \partial\hat{c}_3 \oplus \partial\hat{c}_4} \\ \tau_5(\partial p_7) &= \tau_5(\partial\hat{c}_1 \oplus \partial\hat{c}_2 \oplus \partial\hat{c}_3 \oplus \partial\hat{c}_4 \oplus \partial\hat{c}_5 \oplus \partial\hat{c}_6 \oplus \partial\hat{c}_7) \\ &= \boxed{\partial c_5 \oplus \partial c_4 \oplus \partial c_3 \oplus \partial c_2 \oplus \partial c_1} \oplus \partial\hat{c}_6 \oplus \partial\hat{c}_7. \end{aligned}$$

Previously, we noted that all the partial sums of ∂p_n (as well as ∂q_n) are valid subset sums. It turns out that applying a reflection transform preserves the validity of all partial sums of the new chain.

PROPOSITION 4.22. *Let $\partial\psi$ be a given chain with n elemental chains and an active region $[i, j] \subseteq [1, n]$. If all the partial sums of $\partial\psi$ are valid subset sums, then for any $k \in [j + 1 - i]$, all the partial sums of $\tau_k(\partial\psi)$ are also valid subset sums.*

The proof is straight forward. Suppose the active region of the given chain $\partial\psi$ corresponded to ∂p_k . After a τ_k transform, this portion of the chain will correspond to ∂q_k . Since the low k bits of the indices (binary vectors) of this part of the curve are complements of those before the τ_k transform, it follows that they are valid subset sums too.

The above result enables us to generate new chains by repeated application of τ_k transforms (with decreasing indices) to generate new chains with the guarantee that all partial sums being valid subset sums.

4.3 $p_k \rightarrow \{q_i\}_{i=1}^k$ and $q_k \rightarrow \{p_i\}_{i=1}^k$ structures

The τ transform defined above enables us to generate the whole family of curves $\{q_i\}_{i=1}^k$ from a single curve p_k , and the whole family of curves $\{p_i\}_{i=1}^k$ from q_k . For example, consider the p_n curve with associated chain ∂p_n . Since this chain has the character set $[1, n]$ and contains all the subchains $\partial p_1, \partial p_2, \dots, \partial p_{n-1}, \partial p_n$, we can apply τ_i on this chain to get

$$\begin{aligned}\tau_i(\partial p_n) &= \rho_i(\lambda_i(\partial p_n)) = \partial q_i \oplus \partial \hat{c}_{i+1} \oplus \partial \hat{c}_{i+2} \oplus \dots \oplus \partial \hat{c}_n, \\ &= \partial q_i \oplus \partial p_n[i + 1, n]\end{aligned}$$

for all $i \in [n]$. The last expression means that the last $n - i$ elemental chains are fixed as in ∂p_n , but the first i elemental chains have been reversed to form ∂q_i with new active region $[1, i]$.

For the main example, the effect of the family of transforms $\{\tau\}_{k=4}^9$ on the p_9 curve is shown in Figure 3. Notice that different parts of the curve have differing number of local characters. The lower part of the p_9 curve has more local characters, than the higher part.

Similarly, we can apply τ_i transform on the chain ∂q_n to get

$$\begin{aligned}\tau_i(\partial q_n) &= \rho_i(\lambda_i(\partial q_n)) = \partial c_n \oplus \partial c_{n-1} \oplus \dots \oplus \partial c_{i+1} \oplus \partial p_i, \\ &= \partial q_n[1, n - i] \oplus \partial p_i\end{aligned}$$

for all $i \in [n]$. The last expression means that the first $n - i$ elemental chains are fixed as in ∂q_n , and the last i elemental chains have been reversed to form ∂p_i with new active region $[n + 1 - i, n]$. The effect of the family of transforms $\{\tau\}_{k=4}^9$ on the q_9 curve is shown in Figure 4. In this case, the higher part of the q_9 curve has more local characters.

4.4 Nested τ operations

Here we show an example with two τ transforms applied one after other. Consider the chain ∂q_9 and the nested operations $\tau_5(\tau_8(\partial q_9))$ using main example. Following the




Figure 3: Illustration of the application of $\{\tau\}_{k=4}^9$ on the p_9 curve of main example. The p_9 curve is shown in blue.




Figure 4: Illustration of the application of $\{\tau\}_{k=4}^9$ on the q_9 curve of main example. The q_9 curve is shown in cyan.

above discussion, applying the τ_8 operation first we get

$$\tau_8(\partial q_9) = \partial q_9[1, 1] \oplus \partial p_8.$$

Now applying the τ_5 operation on the result we get

$$\begin{aligned} \tau_5(\tau_8(\partial q_9)) &= \tau_5(\partial q_9[1, 1] \oplus \partial p_8), \\ &= \partial q_9[1, 1] \oplus \partial q_5 \oplus \partial p_8[6, 8]. \end{aligned}$$

It can be noticed that the τ transform operates only a complete chain ∂p_k or ∂q_k for some $k \in [n]$ and never on a partial chain.




Figure 5: Illustration of the application of two τ operations in sequence on the q_9 curve of main example.

5 Generation of all NDPs

The above discussion suggests a natural generalization. By considering either ∂p_n (or ∂q_n) and repeatedly applying the τ transforms we can generate all possible chains. This process can be conveniently represented in the form of a layered directed acyclic graph denoted by $\mathcal{H}$ and is defined next.

5.1 A transformation graph $\mathcal{H}$

DEFINITION 5.1. ($\mathcal{H}$ graph) Let $\mathcal{H} = (V(\mathcal{H}), E(\mathcal{H}))$ be a layered directed graph with vertex set $V(\mathcal{H})$ consisting of chains ∂p_i and ∂q_i for $i \in [n]$, arc set $E(\mathcal{H})$ consisting

the set of transformations τ_j with $j \in [n]$ between the curves. There is a root node representing the chain ∂p_n (or ∂q_n), and the graph $\mathcal{H}$ has $n + 1$ layers.

Since only the active region of a chain can undergo a τ transform, the nodes of the graph represent *only* the active regions, and the arcs represent τ transforms. First, we show the graph with the alternating λ and ρ transforms for clarity, and later we will show the graph using only the composite τ transforms. The root node is the starting chain ∂p_n as shown in Figure 6.

Some remarks on the graph are given:

- (i) Except for the first layer, where we allow $\lambda_n(\partial p_n)$, in all other layers of *local character transformation*, the resulting chain (active region) has a lower numbered local character. The active region monotonically decreases as we go down the graph.
- (ii) The nodes represent the active part of the chain only; the rest of the chain is fixed according to the path traversed from the root. Note that all the chains have a length equal to $n(n + 1)/2$. Since each chain corresponds to a curve, we have as many curves as the number of paths in the graph.
- (iii) Once the local character reaches 1, we cannot go down any further. Actually, we don't even need to go below 3 because since $\partial p_2 = \partial q_2$, and $\partial p_1 = \partial q_1$.
- (iv) The root node is the chain ∂p_n , but we could start with ∂q_n as well. The graph with both λ and ρ transforms has $2n + 1$ levels including the root level.

Since chains and curves are essentially equivalent objects, the same graph structure carries over to curves. For the sake of clarity we again show the graph in Figure 7, with curves instead of chains but with τ transforms which reduce the number of layers from $2n + 1$ to $n + 1$.

5.2 Counting paths in $\mathcal{H}$

The graph $\mathcal{H}$ can be conveniently represented by an upper triangular matrix of size $n \times n$ with elements being the curves p_i/q_i for $i \in [n]$. The initial curve p_n can be thought of in location $(0, 0)$. The first row consists of curves $q_n, q_{n-1}, \dots, q_1$ comprising Q_n . The second row consists of $p_{n-1}, p_{n-2}, \dots, p_1$ comprising P_{n-1} , and so on. The final row consists of either $p_1 = P_1$ or $q_1 = Q_1$, depending on n being even or odd respectively. The $\mathcal{H}$ graph in the form of an upper triangular matrix in the general case is shown below. In each cell of the matrix, the number of incoming paths from p_n , which is a binomial coefficient (proof is given below) is shown in red. The last column shows the total number of paths into each row.




Figure 6: Chains resulting from ∂p_n by repeated alternating λ and ρ transforms. Each node represents only the active region of the chain. This directed acyclic graph has $2n + 1$ levels.

p_n^1	1	2	3	4	5	...	$n-3$	$n-2$	$n-1$	n	#paths
1	$q_n^{\binom{0}{0}}$	$q_{n-1}^{\binom{1}{0}}$	$q_{n-2}^{\binom{2}{0}}$	$q_{n-3}^{\binom{3}{0}}$	$q_{n-4}^{\binom{4}{0}}$	...	$q_4^{\binom{n-4}{0}}$	$q_3^{\binom{n-3}{0}}$	$q_2^{\binom{n-2}{0}}$	$q_1^{\binom{n-1}{0}}$	$\binom{n}{1}$
2		$p_{n-1}^{\binom{1}{1}}$	$p_{n-2}^{\binom{2}{1}}$	$p_{n-3}^{\binom{3}{1}}$	$p_{n-4}^{\binom{4}{1}}$	...	$p_4^{\binom{n-4}{1}}$	$p_3^{\binom{n-3}{1}}$	$p_2^{\binom{n-2}{1}}$	$p_1^{\binom{n-1}{1}}$	$\binom{n}{2}$
3			$q_{n-2}^{\binom{2}{2}}$	$q_{n-3}^{\binom{3}{2}}$	$q_{n-4}^{\binom{4}{2}}$	...	$q_4^{\binom{n-4}{2}}$	$q_3^{\binom{n-3}{2}}$	$q_2^{\binom{n-2}{2}}$	$q_1^{\binom{n-1}{2}}$	$\binom{n}{3}$
4				$p_{n-3}^{\binom{3}{3}}$	$p_{n-4}^{\binom{4}{3}}$	...	$p_4^{\binom{n-4}{3}}$	$p_3^{\binom{n-3}{3}}$	$p_2^{\binom{n-2}{3}}$	$p_1^{\binom{n-1}{3}}$	$\binom{n}{4}$
5					$q_{n-4}^{\binom{4}{4}}$	...	$q_4^{\binom{n-4}{4}}$	$q_3^{\binom{n-3}{4}}$	$q_2^{\binom{n-2}{4}}$	$q_1^{\binom{n-1}{4}}$	$\binom{n}{5}$
$\vdots$						$\ddots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$n-3$							$q_4^{\binom{n-4}{n-4}}$	$q_3^{\binom{n-3}{n-4}}$	$q_2^{\binom{n-2}{n-4}}$	$q_1^{\binom{n-1}{n-4}}$	$\binom{n}{n-3}$
$n-2$								$p_3^{\binom{n-3}{n-3}}$	$p_2^{\binom{n-2}{n-3}}$	$p_1^{\binom{n-1}{n-3}}$	$\binom{n}{n-2}$
$n-1$									$q_2^{\binom{n-2}{n-2}}$	$q_1^{\binom{n-1}{n-2}}$	$\binom{n}{n-1}$
n										$p_1^{\binom{n-1}{n-1}}$	$\binom{n}{n}$

Each element of the first row has a single transformation arc from p_n . In other words $q_{n-i} = \tau_{n+1-i}(p_n)$, for $i = 1, \dots, n$.

DEFINITION 5.2. (*paths*) For an element $(i, j) \in [n] \times [n]$ with $i \leq j$, the number of paths into it is denoted by $\wp(i, j)$, and it equals the sum of all paths into its parent curves in row $i-1$. Formally

$$\wp(i, j) = \sum_{k=i-1}^{j-1} \wp(i-1, k). \quad (5.1)$$

We define $\wp(0, 0) = 1$, and $\wp(0, j) = 0$ for all $j \in [n]$.

From the definition, we have $\wp(1, j) = 1$ for all $j \in [n]$ since all of these curves have only one parent curve p_n . For the second row, we have

$$\begin{aligned} \wp(2, 2) &= \wp(1, 1) = 1, \\ \wp(2, 3) &= \wp(1, 1) + \wp(1, 2) = 2, \\ \wp(2, j) &= \wp(1, 1) + \wp(1, 2) + \dots + \wp(1, j-1) = j-1, \text{ for all } j = 3, \dots, n. \end{aligned}$$

We use the following well known identity (sometimes called as the hockey-stick identity) in the proofs below.

IDENTITY 5.3. For any positive integer $r \in [n]$, we have

$$\binom{n}{r} = \binom{n-1}{r-1} + \binom{n-2}{r-1} + \dots + \binom{r-1}{r-1}.$$

This identity can be easily proved by repeatedly applying Pascal's identity: $\binom{n}{r} = \binom{n-1}{r-1} + \binom{n-1}{r}$.

PROPOSITION 5.4. *The number of paths into an element (i, j) with $i \leq j$ is*

$$\wp(i, j) = \sum_{k=i-1}^{j-1} \wp(i-1, k) = \binom{j-1}{i-1}.$$

Proof. We give the proof by induction. It is true for the first row trivially since $\binom{j}{0} = 1$ for all $j \in [n]$. By induction hypothesis, suppose it is true for row i for some $i \in [n]$. Now consider row $i+1$ and a column j . The parent curve indices for the element $(i+1, j)$ are $(i, i), (i, i+1), \dots, (i, j-1)$. By definition of paths,

$$\begin{aligned} \wp(i+1, j) &= \wp(i, i) + \wp(i, i+1) + \dots + \wp(i, j-1), \\ &= \binom{i-1}{i-1} + \binom{i}{i-1} + \dots + \binom{j-2}{i-1}, \\ &= \binom{j-1}{i}, \text{ (by IDENTITY 5.3)} \end{aligned}$$

and the claim follows. $\square$

DEFINITION 5.5. *The total number of paths into level $r \in [n]$ in $\mathcal{H}$ is the sum of all paths into elements of row r . We denote this quantity by $\beta(r)$.*

THEOREM 5.6. *The number of paths $\beta(r)$ into level $r \in [n]$ follows a Binomial distribution.*

Proof. From the definition above, we have

$$\begin{aligned} \beta(r) &= \sum_{j=r}^n \wp(r, j), \\ &= \binom{r-1}{r-1} + \binom{r}{r-1} + \binom{r+1}{r-1} \dots + \binom{n-1}{r-1}, \\ &= \binom{n}{r}, \text{ (by IDENTITY 5.3),} \end{aligned}$$

and the claim follows. $\square$

PROPOSITION 5.7. *Let $\psi_0 = p_n$. Then the total number of paths into all levels $r \in [n]$ is 2^n , and the maximum number of distinct descendant paths of ψ_0 including ψ_0 itself is $2^n/8$.*

Proof. Since the number of paths into level r is $\binom{n}{r}$, the total number of paths into all levels is easily seen to be 2^n . However, some of these NDPs are not distinct since $\partial p_1 = \partial q_1$ and $\partial p_2 = \partial q_2$. Furthermore, since the pair (p_4, q_4) covers S_4 , it follows that we can skip ∂p_3 and ∂q_3 as well. Thus we can remove the last three columns and last three rows in the matrix form of $\mathcal{H}$ to obtain $2^n/8$ distinct NDPs. $\square$

5.3 A covering by NDPs

Starting with the p_n curve (or q_n curve), we have seen that many new curves can be generated by applying the τ transforms repeatedly. A natural question arises: does this collection of curves cover the solution space S_n ? Next we prove this in the affirmative.

THEOREM 5.8. *The complete set of paths obtained from ψ_0 cover the point set S_n .*

Proof. We give the proof by induction.

(Base Step) Consider the base case $n = 4$. Set $\psi_0 = p_4$ with *active region* being the whole curve. We have only one possibility for descendant paths, namely $\psi_1 = \tau_4(\psi_0)$ which is a q_4 curve. Since the set $\{p_4, q_4\}$ exhausts all the 16 points of the solution space, the hypothesis is true.

(Inductive Step) Suppose the hypothesis is true for $n = k$ where $k > 4$. In other words, if $\psi_0 = p_k$, then the set of all descendant paths of ψ_0 cover S_4 . When we go from k to $k + 1$ we get two clusters S_k of point sets; one located at 0 and one at c_{k+1} . We have $\psi_0 = p_{k+1}$ with active region being whole of p_{k+1} . By performing $\tau_{k+1}(\psi_0)$, we obtain q_{k+1} . Now applying $\tau_k(q_{k+1})$, we obtain p_k which along with q_{k+1} (locally q_k) encloses the second cluster S_k located at c_{k+1} . All points within this cluster S_k can be reached by hypothesis.

Now we need to show that all points in the first cluster can be reached from ψ_0 as well. This is easy since p_{k+1} has multiple local characters of $\{p_4, p_5, \dots, p_k, p_{k+1}\}$. By choosing the local character p_k , i.e., by applying $\tau_k(p_{k+1})$ we obtain the q_k curve in the first cluster located at 0. This q_k curve along with p_{k+1} (locally p_k) encloses all the points of the first cluster. By hypothesis all points in this cluster S_k can be reached. Hence it is true that all points in S_n are covered by paths obtained from ψ_0 . $\square$

DEFINITION 5.9. *Let Ψ denote the set of all descendant paths, obtained from ψ_0 , including ψ_0 itself.*

Recall that the NDPs are constructed to be independent of the magnitude of elements in $\mathbf{a}$. Thus the family of paths Ψ will remain non-decreasing for all instances of size n . Thus, we have proved our first claim:

THEOREM 5.10. *The set of non-decreasing paths Ψ is a universal geometric structure for all instances of size n .*

The complete set of paths, obtained starting from p_n (or equivalently from q_n), for the *main example* are shown in Figure 8. The point set that would appear like a random set of points in the Cartesian plane without the paths, once we overlay the paths suddenly reveals a hidden order and symmetry. The fascinating symmetries of paths and subspaces, as well as coverage of the point set is evident in the figure. While the configuration of points will change with each instance, the family Ψ will track them. All the $2^{9-3} = 64$ paths between 0 and C_9 can be seen (in fact be counted visually). The




Figure 8: Illustration of the set of all NDPs Ψ for the main example. The p -curves are in blue, and the q -curves are in cyan. The symmetries of this family of curves and the subspaces are clearly evident. The dotted black line is the OL $y = T$, while the dotted red line is the complement $y = A_n - T$. The highlighted circle corresponds to the solution with target value T , while the highlighted square corresponds to the complement, i.e., $A_n - T$. After all, the solution space of the SSP has a rich and beautiful structure!

path containing the solution (marked by a circle) corresponds to the following sequence of transformations:

$$\begin{aligned} \partial\psi_{sol} &= \tau_5(\tau_7(\tau_8(\partial q_9))), \quad (\text{if starting from } \partial q_9) \\ \text{or } \partial\psi_{sol} &= \tau_5(\tau_7(\tau_8(\tau_9(\partial p_9)))) \quad (\text{if starting from } \partial p_9). \end{aligned}$$

5.4 Segments and Path multiplicity

DEFINITION 5.11. (*segment*) Let $s_i = [z_i^-, z_i^+]$ be an interval with $z_i^-, z_i^+ \in S_n$. We call s_i a segment if and only if it is on at least one NDP $\psi \in \Psi$.

We have $|\Psi| = \frac{2^n}{8}$ distinct NDPs in the geometry, and each NDP has $\frac{n(n+1)}{2}$ edges. Thus there are $\frac{2^n n(n+1)}{16}$ total segments in the geometry. It can be proved that the number of unique segments is only $2^n + 2$ for all $n > 3$, and that for $n = 3$ it is 2^3 .

DEFINITION 5.12. (*path multiplicity*) To each unique segment $s_i, i \in [2^n + 2]$ in the geometry, we associate a positive integer ν_i , which is the number of NDPs that go through s_i . We call ν_i the path multiplicity of segment s_i .

For example, for the first segment $s_1 = [0, b_1 + \iota a_1]$, $\nu_1 = \frac{2^n}{8}$, since all the NDPs go through it. Similarly the last segment with end points $[(B_n - b_1) + \iota(A_n - a_1), B_n + \iota A_n]$ also has a path multiplicity of $\frac{2^n}{8}$. However, the segment $[b_2 + \iota a_2, b_1 + b_2 + \iota(a_1 + a_2)]$ has only one NDP going through it. In general, the path multiplicity of an arbitrary segment in an NDP $\psi \in \Psi$ lies in the interval $[1, \frac{2^n}{8}]$.

Next, consider the segments of Figure 8 that intersect the OL. With not too much difficulty, one can see that there are 25 distinct segments that intersect the OL with multiplicities (16, 8, 4, 2, 2, 1, 1, 8, 4, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 2) from left to right respectively. Also note that the sum of these multiplicities is equal to $64 = 2^{9-3}$, the total number of distinct NDPs in the *main example*. The first intersection point is that of p_7 which contains all the subspaces leading to a multiplicity of 16. This same pattern continues for larger n with many intersection points having an exponential multiplicity. The main thing to note is that many of $2^n/8$ curves do not directly intersect the OL; they only do through higher curves (subspaces) in which they are contained. When all the elements of $\mathbf{a}$ are distinct, there is a gradual rise of the graph as we go from left to right. This staggering causes many intersection points to have high multiplicities. We will refer to this discussion again in Section 9.

5.5 A transformation vector

We have shown that the set of NDPs Ψ derived from ψ_0 covers the solution space.

Given a binary vector $\mathbf{x}$ what is the NDP ψ that contains it? The answer is given by the following theorem.

THEOREM 5.13. Given any binary vector $\mathbf{x} \in \{0, 1\}^n$, and $\psi_0 = q_n$ (or p_n), there is at least one NDP $\psi = K(\psi_0)$ that contains it, where $K = (k_l, k_{l-1}, \dots, k_2, k_1)$ is the sequence of transformations τ_{k_i} on ψ_0 . Furthermore, the sequence of transformations

encoded as a binary vector K , that produce ψ from ψ_0 are encoded in the vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$ itself as

$$\text{index}(K) = \sum_{i=1}^n |x_{i+1} - x_i| b_i,$$

where

$$x_{n+1} = \begin{cases} 1 & \text{if } \psi_0 = q_n \\ 0 & \text{if } \psi_0 = p_n. \end{cases}$$

Proof. Let $\mathbf{x}$ be the given binary vector with coordinates given by $z = (\mathbf{b} \cdot \mathbf{x}, \mathbf{a} \cdot \mathbf{x})$. Recall that the space S_n comprises of two copies of S_{n-1} , one at 0 and the other at c_n . We refer to these as the left (L) and right (R) clusters. If $x_n = 1$ then it is clear that the point z is in the cluster R , whereas if $x_n = 0$ then it belongs to cluster L . We pick the cluster that contains z and descend into S_{n-1} . Then we check if x_{n-1} is 1 or 0 and descend into the appropriate cluster S_{n-2} . We keep repeating this following a binary tree until x_1 at which point we have located the point. If $\psi_0 = q_n$ and $x_n = 1$, then the point is in the cluster R we don't need to do anything. On the other hand if $x_n = 0$ we need to go to cluster L , which requires the transformation $\partial p_n = \tau_n(\partial q_n)$. It is straight forward to see that a τ_k operation is needed whenever $x_{k+1} = 1$ and $x_k = 0$, or $x_{k+1} = 0$ and $x_k = 1$. Alternatively, we need a τ_k transform when there is a change in consecutive bits $1 \rightarrow 0$ or $0 \rightarrow 1$. When a τ_k operation is used we set the k th bit of $\text{index}(K)$ to 1, otherwise set it to 0. Doing this we obtain the formula for $\text{index}(K)$ and the claim follows. $\square$

EXAMPLE 5.14. Consider the solution vector $\mathbf{x} = (1, 0, 1, 1, 0, 1, 1, 0, 1)$ of the main example, with bit positions $9, 8, \dots, 1$ respectively. Then using the above theorem, we can compute the index of the K vector as

$$\text{index}(K) = b_8 + b_7 + b_5 + b_4 + b_2 + b_1 = 219$$

with the sequence of τ transforms given by

$$\partial\psi = \tau_1(\tau_2(\tau_4(\tau_5(\tau_7(\tau_8(\partial q_9)))))).$$

Note that in this example, we can in fact stop at τ_5 since the point corresponding to the solution vector becomes part of a fixed chain after this.

THEOREM 5.15. The Subset Sum Problem has a solution, if and only if, at least one path $\psi \in \Psi$, intersects the orbital line at a vertex point.

How do we find all the points on the OL? A brute force way is to exhaustively search all the paths in Ψ , where for each path we check all the points of each curve if any equals the target value. Clearly this is infeasible!

By understanding the relationship between the NDPs in Ψ and using that relationship, we can indeed obtain a better algorithm. This will need a few more ideas which are covered in the subsequent sections.

5.6 Wormhole paths of $\mathcal{H}$

We have seen that for any path in $\mathcal{H}$, the active region is a monotonically decreasing function of path length (number of transformations). In this section we examine the lower and upper boundaries of the active region at any point in the path and how they create *wormhole* paths.

Consider the chain is ∂p_n . If we apply the τ_j operation on it where $j < n$ we get

$$\tau_j(\partial p_n) = \rho_j(\lambda_j(\partial p_n)) = \partial q_j \oplus \partial p_n[j+1, n].$$

The active region of the initial chain is the interval $[1, n]$ which corresponds to the whole curve $p_n(t)$, i.e., $[(0, 0), (B_n, A_n)]$. After the transformation, the new active region corresponds to $[(0, 0), (B_j, A_j)]$. Now consider a further transformation on ∂q_j given by $\tau_k(\partial q_j \oplus \partial p_n[j+1, n]) = \partial q_j[1, j-k] \oplus \partial p_k \oplus \partial p_n[j+1, n]$. Now the active region of the new curve is $[(B_j - B_k, A_j - A_k), (B_j, A_j)]$. We notice from this example that when we go from a p -curve to a q -curve, the upper end of the active region is lowered, and when we go from a q -curve to a p -curve the lower end point is raised. In either transformation, the length of the active region (measured by y range) decreases. For any path in $\mathcal{H}$, the active region monotonically decreases with number of transformations, by alternatingly lowering the upper end point or raising the lower end point.

EXAMPLE 5.16. *In the main example, we have seen previously that the curve containing the solution point can be obtained from p_9 via the sequence of transformations: $(\tau_9, \tau_8, \tau_7, \tau_5, \tau_4)$. In Figure 9 we show the intervals (y coordinate only) containing the curves after each successive τ transformation. The final curve is that of q_4 after transformation τ_4 . The OL is shown as a blue line. We notice that the y -interval $[395296, 748620] \in p_9$ corresponds to the projection of q_4 onto p_9 .*

DEFINITION 5.17. *For an arbitrary path π of length r in $\mathcal{H}$, represented by*

$$\pi = (k_1, k_2, \dots, k_r)$$

where $k_1 > k_2 > \dots > k_r$ and $k_i \in [n]$ for all $i \in [r]$, let $\ell := (\ell_1, \ell_2, \dots, \ell_r)$ be the sequence of y values for the lower curve, and let $\bar{h} := (\bar{h}_1, \bar{h}_2, \dots, \bar{h}_r)$ be the sequence of y values for the upper curve. Note that for any $i \in [r]$, the pair $(\ell_i, \bar{h}_i)$ corresponds to a complete p_j/q_j curve for some $j \in [n]$.

PROPOSITION 5.18. *For a path π of length r in $\mathcal{H}$, represented by*

$$\pi = (k_1, k_2, \dots, k_r)$$

where $k_1 > k_2 > \dots > k_r$ and $k_i \in [n]$ for all $i \in [r]$, we have

$$\ell_i = \begin{cases} \sum_{j=1}^i (-1)^{(j+1)} A_{k_j} & \text{if } i \text{ is even,} \\ \sum_{j=1}^{i-1} (-1)^{(j+1)} A_{k_j} & \text{if } i \text{ is odd,} \end{cases}$$

where the empty sum is treated as zero, and

$$\bar{h}_i = \alpha_i + A_{k_i}.$$




Figure 9: Illustration of the path $\pi = (\tau_9, \tau_8, \tau_7, \tau_5, \tau_4)$ for the main example. The top and bottom lines show the limits of the curve after each transformation. The OL is shown in blue.

The proof is omitted as it is simply a consequence of the transformation of chains. The sequences ℓ and $\hbar$ given by the above proposition have the form

$$\begin{aligned} \ell &= (0, (A_{k_1} - A_{k_2}), (A_{k_1} - A_{k_2}), (A_{k_1} - A_{k_2}) + (A_{k_3} - A_{k_4}), \\ &\quad (A_{k_1} - A_{k_2}) + (A_{k_3} - A_{k_4}), \dots), \\ \hbar &= (A_{k_1}, (A_{k_1} - A_{k_2}) + A_{k_2}, (A_{k_1} - A_{k_2}) + A_{k_3}, (A_{k_1} - A_{k_2}) + (A_{k_3} - A_{k_4}) + A_{k_4}, \\ &\quad (A_{k_1} - A_{k_2}) + (A_{k_3} - A_{k_4}) + A_{k_5}, \dots) \end{aligned}$$

as can be verified. It can be seen that ℓ is alternately staying the same and increasing, while $\hbar$ is alternately decreasing and staying the same.

DEFINITION 5.19. (wormhole) For any transformation path $\pi \in \mathcal{H}$, let ℓ^π and $\hbar^\pi$ be the associated lower and upper curves respectively. We then call the region enclosed by the curves ℓ^π and $\hbar^\pi$ as a wormhole path or simply a wormhole due to its shape.

DEFINITION 5.20. (Valid Wormhole for a given OL) Given any path π from p_n in $\mathcal{H}$, and a given T we say that the path π is valid if and only if the associated wormhole contains T . Alternatively, a path π of length k is valid if and only if

$$\ell_i^\pi \leq T \leq \hbar_i^\pi,$$

for all $i \in [k]$.

Next we give a larger example illustrating the fact that for a given OL, not every path in $\mathcal{H}$ need contain it.

EXAMPLE 5.21. *This is a random example with $n = 30, m = 40$, with some path $(30, 27, 25, 24, 21, 17, 16, 12, 9, 7, 5)$ in $\mathcal{H}$. The corresponding intervals are shown along with the OL in blue.*




Figure 10: Illustration of an arbitrary path to level $r = 20$ for a $n = 30$ example. The OL (blue line) is contained in the initial part of the path but later misses it. This shows that not every path in $\mathcal{H}$ contains a given OL $y = T$.

5.7 Number of OLs into each curve/level for a given $T \in p_n$

We have seen above that not all wormholes associated with the paths in $\mathcal{H}$ contain a given OL $y = T$. To provide a clear picture for the upcoming analysis in the subsequent sections, we provide some definitions below:

DEFINITION 5.22.

Let $\beta(r) := \# \text{ transformation paths into level } r \text{ in } \mathcal{H}$,

$\beta(r, i) := \# \text{ transformation paths into curve } p_i/q_i \text{ in level } r \text{ in } \mathcal{H}$,

$\beta_T(r, i) := \# \text{ valid OLs into curve } p_i/q_i \text{ in level } r \text{ in } \mathcal{H} \text{ for a given } T \in p_n$,

$\beta_T(r) := \# \text{ valid OLs into all curves of level } r \text{ in } \mathcal{H} \text{ for a given } T \in p_n$,

$\hat{\beta}_T(r, i) := \# \text{ distinct valid OLs into curve } p_i/q_i \text{ in level } r \text{ in } \mathcal{H} \text{ for a given } T \in p_n$,

and $\hat{\beta}_T(r) := \# \text{ distinct valid OLs into all curves of level } r \text{ in } \mathcal{H} \text{ for a given } T \in p_n$.

We have proved earlier that $\beta(r, i) = \binom{n+1-i}{r-1}$ and $\beta(r) = \binom{n}{r}$.

The following proposition follows directly from the above definition.

PROPOSITION 5.23. *For any given instance $(\mathbf{a}, T)$ it is true that*

$$\hat{\beta}_T(r, i) \leq \beta_T(r, i) \leq \beta(r, i) = \binom{n+1-i}{r-1}$$

and

$$\hat{\beta}_T(r) \leq \beta_T(r) \leq \beta(r) = \binom{n}{r}.$$

6 Intersection Points and Local Coordinates

We have seen that there are at most $2^n/8$ distinct NDPs from 0 to C_n . For a given orbital line $y = T$, all these paths intersect the OL. We are interested in understanding the relation between these intersection points and this section is devoted for this task.

6.1 Edges and Local Reference Frames

Let ψ be an NDP with associated chain $\partial\psi$. The chain $\partial\psi$ comprises of $\frac{n(n+1)}{2}$ links, and the corresponding curve ψ comprises of the same number of line segments which are referred to as *edges* (formally defined below). Recall that the points on the curve ψ are given by

$$\psi(i) = \sigma(\partial\psi, i), \quad \text{for all } i \in \left[\frac{n(n+1)}{2}\right]$$

where $\psi(0) = 0$.

DEFINITION 6.1. (*Edge*) Let ψ be a NDP. The line segments of ψ are called edges and denoted by e_i , where $i \in \left[\frac{n(n+1)}{2}\right]$. The lower end point of edge e_i is denoted by z_i^- and the upper end point by z_i^+ . Thus,

$$\begin{aligned} z_i^- &= x_i^- + \iota y_i^- = \sigma(\partial\psi, i-1), \\ \text{and } z_i^+ &= x_i^+ + \iota y_i^+ = \sigma(\partial\psi, i). \end{aligned}$$

We associate the half open interval $[z_i^-, z_i^+)$ to each edge e_i . Note that the point z_i^+ itself does not belong to e_i , to avoid double counting of vertex points. Since each edge corresponds to a link in the chain, for an edge e_i we have an associated link d_j for some $j \in [n]$. We denote the length of the edge e_i by $|e_i|$ and define it (using only the y coordinates) as

$$|e_i| = y_i^+ - y_i^- = \Im(d_j).$$

While the Euclidean length of e_i is strictly positive (since $\Re(d_j) > 0$ and $\Im(d_j) \geq 0$), the length $|e_i|$ however can be zero if $a_j = a_{j-1}$.

Comparing the above definition to Definition 4.1, one can see that an *edge* is nothing but a translated *link*. Another way of saying this is that a *chain* is a sequence of *links*, whereas a NDP is a sequence of *edges*.

DEFINITION 6.2. (*Local Reference Frame of an edge e_i*) Let e_i be an edge on a NDP ψ , with end points z_i^- and z_i^+ corresponding to a difference d_j for some $j \in [n]$. We imagine a Cartesian coordinate system located at z_i^- , and call this the local reference frame (LRF) of the edge e_i . We call the point z_i^- as the origin of the edge e_i . An arbitrary point $z \in \mathbb{C}$ has a local value z_i^{loc} in the LRF of e_i , given by

$$z_i^{loc} = z - z_i^-.$$

The point z in global coordinates is simply $z = z_i^- + z_i^{loc}$.

DEFINITION 6.3. (*Valid points on an edge*) A point z is said to be valid point of edge e_i , if it lies in between the end points of the edge. For a valid point z , there exists a rational $\delta \in [0, 1)$, such that $z = z_i^- + \delta d_j$. Then the coordinates of z in the LRF of the edge e_i , denoted z_i^{loc} , are

$$z_i^{loc} = x_i^{loc} + \iota y_i^{loc} = \delta d_j.$$

For a valid point z in the LRF of e_i , we simply write $z \in e_i$.

6.2 Local coordinates of intersection points

Next we look at the intersection point of a NDP ψ with a horizontal line $y = y_0$.

PROPOSITION 6.4. Given a positive integer $y_0 \in [0, A_n]$, the number of points belonging to NDP ψ , that have the y -coordinate equal to y_0 , is either 1 or ∞ .

Proof. If all the elements of $\mathbf{a}$ are distinct, then it follows that $\Im(d_j) > 0$ for all $j \in [n]$. Since $0 < y_0 \leq A_n$, there is a unique edge $e_i \in \psi$, such that $y_i^- \leq y_0 < y_i^+$, and there is a unique intersection point of line $y = y_0$ with the edge e_i . In the case one or more elements of D have $\Im(d_j) = 0$, it is possible for one or more edges e_i to coincide with the line $y = y_0$. In this case, we have infinite number of points belonging to ψ that have y coordinate equal to y_0 . $\square$

For the case when all elements of $\mathbf{a}$ are distinct, let the intersection point of ψ with $y = y_0$ line be denoted $z_\psi = x_\psi + \iota y_\psi$ where $y_\psi = y_0$. Furthermore, let $e_i \in \psi$ (with associated link d_j) contain the intersection point z_ψ . From the above definitions, we can write $z_\psi = z_i^- + z_i^{loc}$ where $z_i^{loc} = \delta_\psi d_j$ and $\delta_\psi = \frac{y_i^{loc}}{|e_j|}$.

For the case where the edge e_i coincides with the $y = y_0$ line, we use the following prescription for local coordinates where we set

$$\delta_\psi := \begin{cases} 0 & \text{if } |e_j| = 0 \\ \frac{y_\psi^{loc}}{|e_j|} & \text{otherwise.} \end{cases} \quad (6.1)$$

With this, the local x_i^{loc} coordinate is unambiguously determined as

$$x_i^{loc} = \delta_\psi \Re(d_j).$$

Thus, eventhough there are infinite number of intersection points (global coordinates) we have only a single local coordinate!

REMARK 6.5. *If there are h edges of ψ , that are all horizontal and next to each other, and coincide with the $y = y_0$ line, then we would get h local coordinates (one for each edge) rather than an infinite number of them using the above definition for local coordinates.*

6.3 A Bit Of Relativity

Consider a pair of complementary curves p_k and q_k , and a horizontal line $y = y_0$ where $y_0 \in [0, A_k]$. Since both p_k and q_k curves start at 0 and end at C_k , both curves will intersect the $y = y_0$ line. Let the intersection points be denoted $z_p = x_p + \iota y_p$ and $z_q = x_q + \iota y_q$ respectively. It is clear that $y_p = y_0 = y_q$.

These points are in *global* coordinates (relative to the origin of subspace translate S_k that contains the p_k and q_k curves). For simplicity we assume that each intersection point lies in a unique edge. (If multiple edges intersect the line, we will show how it can be handled easily later on). Let z_p belong to some edge $e_j \in p_k$ with lower end point at z_j^- . Similarly, let z_q belong to some edge $e_i \in q_k$, with lower vertex at z_i^- . Using the LRF of each edge we can write this as

$$\begin{aligned} z_q^{loc} &= z_q - z_i^- \\ z_p^{loc} &= z_p - z_j^- = (z_q - z_i^-) + (z_p - z_q) + (z_i^- - z_j^-). \end{aligned}$$

Thus

$$z_p^{loc} = z_q^{loc} + (z_p - z_q) + (z_i^- - z_j^-). \quad (6.2)$$

Since both z_i^- and z_j^- are coordinates relative to the origin of space S_k , the translation from z_i^- to z_j^- is simply $z_j^- - z_i^-$. We denote this translation by $w_{ij} := z_j^- - z_i^-$, and can view it as a directed arrow from z_i^- to z_j^- .

If we consider only the y coordinate, it can be readily seen that

$$y_p^{loc} = y_q^{loc} + (y_i^- - y_j^-) = y_q^{loc} + \Im(\overline{w}_{ij}) \quad (6.3)$$

where $\overline{w}_{ij}$ is complex conjugate of w_{ij} . This tells us that if we know y_q^{loc} in the q_k curve and the translation w_{ij} we can obtain the local coordinate y_p^{loc} in the p_k curve, and vice versa.

In the general case, we can have multiple edges $\{e_{i_1}, e_{i_2}, \dots, e_{i_r}\}$ where $r < k$, intersecting the OL instead of a single edge e_i of curve q_k . Similarly, we could have $\{e_{j_1}, e_{j_2}, \dots, e_{j_s}\}$ where $s < k$, intersecting the OL instead of a single edge e_j of curve p_k . In view of Remark 6.5, we have r local coordinates on the q_k curve and s local

coordinates on the p_k curve. Why are $r, s < k$? A brief thought reveals that if $r \geq k$ then all the elements of $\{d_1, d_2, \dots, d_k\}$ have imaginary part equal to zero, which further implies that all $a_i = 0, \forall i \in [k]$. Since we assumed that $\mathbf{a}$ consists of positive numbers, this situation cannot happen. The above arguments are summarized in the following proposition.

PROPOSITION 6.6. *Given a point on a edge of p_k (in local coordinates), we can find at most $k-1$ points, in $k-1$ edges of the complementary curve q_k , such that all these points fall on the same OL. Furthermore, all the points on q_k can be obtained by translations from the given point on p_k . The same is true when p_k is replaced by q_k .*

Since the NDPs form a rigid structure, for any fixed edges $e_1 \in p_k$ and $e_2 \in q_k$ in the same subspace, their relative translation remains invariant, as we go from one instance of subspace S_k to another. Since each curve has $k(k+1)/2$ edges, we can compute all the $(k(k+1)/2)^2$ translations. These translations remain the same for any translate of S_k in the point set. Using these translations, given any point in any edge of p_k , we can determine the image point (again in local coordinates) in a suitable edge in q_k , such that the global values (relative to origin of S_k) have the same y coordinate. We will build on this in the next section.

6.4 Interacting edges

DEFINITION 6.7. (*Edge interaction*) Let e_r be an edge of the curve p_k , defined by the interval $[z_r^-, z_r^+]$. Similarly let e_s be an edge of the curve q_k , defined by the interval $[z_s^-, z_s^+]$. We say that e_r and e_s interact, if they have non-empty overlap in the y -coordinate, i.e.,

$$[y_r^-, y_r^+] \cap [y_s^-, y_s^+] \neq \emptyset. \quad (6.4)$$

Otherwise, the edges are said to be non-interacting. We write $e_r \cap e_s \neq \emptyset$, for simplicity, if they are interacting.

In view of this definition, it can be seen that an edge e_r on curve p_k can interact with *more than one* edge of curve q_k and vice versa. For each of the $k(k+1)/2$ edges of p_k , we can determine which edges of q_k interact with it.

DEFINITION 6.8. (*Interaction length*) The length of the interval over which two edges e_r and e_s interact is called the interaction length, which is denoted by $|e_r \cap e_s|$ and is easily seen to be

$$|e_r \cap e_s| = \min(y_r^+, y_s^+) - \max(y_r^-, y_s^-).$$

The edges e_r and e_s are said to be interacting if and only if $|e_r \cap e_s| \geq 0$.

EXAMPLE 6.9. Let $e_p \in p_k$ and $e_q \in q_k$, such that $|e_p \cap e_q| \geq 0$. Then given $z_p^{loc} \in e_p$, we have $y_q^{loc} = y_p^{loc} + \Im(\overline{w}_{pq})$. We have three cases illustrated in Figure 11.

- (i) The local y values are related by $y_q^{loc} = y_p^{loc} + (y_p^- - y_q^-)$. Here, the coordinates in both edges are valid.

- (ii) Given a valid $y_q^{loc} \in e_q$, we have $y_p^{loc} = y_q^{loc} + (y_q^- - y_p^-) < 0$, which is an invalid point.
- (iii) Similarly, given a valid $y_p'' \in e_p$, we have $y_q^{''loc} = y_p^{''loc} + (y_p^- - y_q^-) > |e_q|$, which is above the edge. Again, this is an invalid point.




Figure 11: Illustration of two interacting edges. All coordinates shown are relative to the origin of S_k that contains the curves p_k and q_k .

The above ideas can be summarized by the following proposition.

PROPOSITION 6.10. *Let ψ_1 be a NDP with active region corresponding to p_j for some $j \in [n]$. Let ψ_2 be a derived path of ψ_1 , given by the relation $\partial\psi_2 = \tau_k(\partial\psi_1)$, where $k < j$. Let $e_p \in \psi_1$ and $e_q \in \psi_2$, such that $|e_p \cap e_q| \geq 0$. Then given a point $z_p^{loc} \in e_p$, we have $y_q^{loc} = y_p^{loc} + \Im(\overline{w}_{pq})$.*

If we look at Figure 8, in particular the OL, we see many edges intersecting it as we go from left to right. All these intersection points have local coordinates in the reference frames of corresponding edges. With the work done in previous section and this, we have a way to hop from one edge to another. As the reader may guess, a repeated application of Proposition 6.10, starting with an edge on p_n , can enable us to visit all the intersection points without leaving the OL.

6.5 An unique representation of an edge

In order to hop on the OL from one edge to another, we need a way to refer to the edges, as there are an exponential number of edges. Fortunately, the graph structure in Figure 7 itself provides part of the solution, as the curves are arranged in levels, i.e., equivalence classes. The chain ∂q_j consists of $j(j+1)/2$ links, but only j unique links, i.e., $d_1, d_2, \dots, d_j$, and the links repeat as we go along the chain. Thus, for the curve q_j , the edges which correspond to links, repeat. As a result, we need a way to distinguish two edges of the same kind at different locations on the curve.

To handle this, we provide a unique 5-tuple representation of an edge, defined below.

DEFINITION 6.11. (5-tuple representation) For an edge on a path ψ , we associate a 5-tuple: (level number, curve type, edge number, curve index, instance number) which are defined as

level $\in [n]$: number of τ operations needed on ∂p_n to obtain $\partial\psi$,
 curve type $\in \{p\text{-type}, q\text{-type}\}$: the type of the active region in $\partial\psi$
 edge number $\in [n]$: j if the edge corresponds to d_j
 curve index $\in \{1, \dots, n\}$: j if active region corresponds to p_j or q_j
 instance number $\in [n]$: k if it is the k th copy of edge d_j .

Thus, for each edge e_i we can associate a unique 5-tuple $F(e_i) = (f_1, f_2, f_3, f_4, f_5)$, where f_1 is the level number, f_2 is the curve type, f_3 is the edge number, f_4 is the curve index and f_5 is the instance number. We will show later that the maximum number of unique 5-tuples is $\frac{n(n+1)(n+2)(n+3)}{24}$.

6.6 Number of interacting edge pairs

Consider a pair of complementary curves p_k and q_k for some $k \in [n]$. Since there are $k(k+1)/2$ edges in each curve, the maximum number of pairs is trivially $(\frac{k(k+1)}{2})^2$. However, unless $\mathbf{a} = \mathbf{0}$, not all pairs will be interacting. Here, we provide bounds on the minimum and maximum number of interacting edges between two complementary curves. Recall that two edges are said to be interacting if their interaction length is non-negative, i.e., even if their intersection is a single point.

DEFINITION 6.12. (Number of interacting edges J_k) Let $k \in [n]$ be a positive integer. Then we define $J_k(\mathbf{a})$ to be the number of interacting edge pairs between curves p_k and q_k .

PROPOSITION 6.13. Let $\mathbf{a}_{diss}$ be a dissociated set and let $3 \leq k \leq n$. Then $J_k(\mathbf{a}_{diss}) = k^2 + k - 5$.

Proof. Since $\mathbf{a}_{diss}$ is dissociated, the subset sums of p_k and q_k (except for the 6 common points) will all be distinct. Thus, the number of unique points between both curves is $k(k+1) - 4$. If we use this set of points to interpolate both curves, we get $k(k+1) - 5$ disjoint segments where each segment on q_k interacts with only one segment on p_k . Thus, the total number of interactions between q_k and p_k is $J_k(\mathbf{a}_{diss}) = k^2 + k - 5$. $\square$

To maximize the total number of interacting edge pairs between two complementary curves, with the requirement that $\mathbf{a} \neq \mathbf{0}$, a momentary thought shows that we can take $\mathbf{a} \in CP$ where all elements of $\mathbf{a}$ are equal.

PROPOSITION 6.14. If $\mathbf{a} \in CP$, the case when all elements are equal to a positive integer constant (say) c , then $J_k(\mathbf{a}_{CP}) = \frac{k^3 + 6k^2 - k}{6}$.

Proof. Note that the number of unique sums in the curve p_k (as well as q_k) is $k+1$. Consider the curve p_k and let's go from the first edge, all the way to the $k(k+1)/2$ th edge

one by one. For the first edge of p_k (sum going from 0 to c), the number of interactions is 1. For each edge of p_k leading from sum jc to $(j+1)c$, where $1 \leq j \leq k-1$, the number of edges of q_k that interact with it can be seen to be $k+1-j$ (It will be helpful to consider the Filling Boxes Model of Section 3). Meanwhile, there are $(j+1)$ edges of p_k that have lower vertex sum to be jc , all of which interact with the $k+1-j$ edges of q_k . Therefore, the total number of interactions is

$$\begin{aligned} J_k(\mathbf{a}_{CP}) &= 1 + \sum_{j=1}^{k-1} (j+1)(k+1-j) \\ &= 1 + \frac{k^3 + 6k^2 - k - 6}{6}, \\ &= \frac{k^3 + 6k^2 - k}{6}. \end{aligned}$$

□

COROLLARY 6.15. *For a random nonzero $\mathbf{a}$, and $k \in [n]$, we have*

$$k^2 + k - 5 \leq J_k(\mathbf{a}) \leq \frac{k^3 + 6k^2 - k}{6}.$$

7 The Orbital Graph G_0

In this section, we describe the details of a layered directed acyclic graph, referred to as the *Orbital Graph* that allows us to visit all the intersection points without leaving the OL.

DEFINITION 7.1. *Let $P_k := \{p_1, p_2, \dots, p_k\}$ be the set of the first k curves of p -type, and $Q_k := \{q_1, q_2, \dots, q_k\}$ be the set of the first k curves of q -type. The number of edges in P_k is denoted by $|P_k|$. Similarly for Q_k .*

The starting point for the construction of the layered directed acyclic graph is the graph structure in Figure 7 which can be written compactly as:

$$p_n \xrightarrow{\{\tau\}} Q_n \xrightarrow{\{\tau\}} P_{n-1} \xrightarrow{\{\tau\}} Q_{n-2} \xrightarrow{\{\tau\}} \dots \xrightarrow{\{\tau\}} P_2 / Q_2 \xrightarrow{\{\tau\}} Q_1 / P_1, \quad (7.1)$$

where each arrow represents multiple τ transforms. This graph structure has $n+1$ layers, where each layer corresponds to a family of curves of either q -type or p -type but not both, and they alternate. While the object of study in the graph will be edges (nodes), we can still use the curves as a packaging of edges. Thus, we view each curve p_j (or q_j), not as a curve but a collection of $\frac{j(j+1)}{2}$ disjoint edges(nodes). Note that we will be using this new compact graph going forward.

Also, note that one could use the alternate graph,

$$q_n \xrightarrow{\{\tau\}} P_n \xrightarrow{\{\tau\}} Q_{n-1} \xrightarrow{\{\tau\}} P_{n-2} \xrightarrow{\{\tau\}} \dots \xrightarrow{\{\tau\}} Q_2 / P_2 \xrightarrow{\{\tau\}} P_1 / Q_1,$$

instead where we start with q_n . We can use either form, and we use Equation (7.1).

The number of edges in each layer is given by the following proposition.

PROPOSITION 7.2. *The sets P_k and Q_k have $\frac{k(k+1)(k+2)}{6}$ edges each.*

Proof. We know that the curve p_j has $\frac{j(j+1)}{2}$ edges. Summing over j as it ranges in $[k]$, we have

$$|P_k| = \sum_{j=1}^k \frac{j(j+1)}{2} = \frac{k(k+1)(k+2)}{6}.$$

Similarly for $|Q_k|$. □

Next we estimate the number of arcs between two adjacent levels, say Q_k and P_{k-1} (or equivalently P_k and Q_{k-1}) for some $4 \leq k \leq n$ in the worst case.

PROPOSITION 7.3. *Let $4 \leq k \leq n$. For any given $\mathbf{a}$, the maximum number of interacting edges between all the edges of Q_k and all the edges of P_{k-1} , in the worst case is $\approx \frac{k^5}{120}$.*

Proof. To compute the total number of interacting edges between Q_k and P_{k-1} , we note that for each $q_i \in Q_k$, we have the transformation structure given by

$$q_i \xrightarrow{\{\tau\}_{k=4}^i} \{p_i, p_{i-1}, p_{i-2}, \dots, p_4\}.$$

Since Q_k is composed of curves $(q_k, q_{k-1}, \dots, q_4)$, we have $(k-i)$ copies of transformations $q_i \rightarrow p_i$. Therefore,

$$\begin{aligned} \# \text{ interacting edges} &\leq \sum_{i=4}^{k-1} (k-i) J_i(\mathbf{a}_{CP}) \quad (\text{Using Proposition 6.14}) \\ &= \sum_{i=4}^{k-1} (k-i) \left(\frac{i^3 + 6i^2 - i}{6} \right) \\ &= \frac{k^5 + 10k^4 - 5k^3 - 10k^2 - 2276k + 6000}{120} \\ &\approx \frac{k^5}{120} \text{ for large } k. \end{aligned}$$

□

7.1 Determining the root node e_0

To determine the root node, we search the edges of curve p_n for the edge that contains the given target value T . Since p_n is a NDP, and we have a formula for the indices, we could use binary search to locate the edge(s) containing the intersection point with $y = T$. This can be done in logarithmic time. We have two possible cases as shown by Proposition 6.4:

- (1) Only one intersection point; in this case we can identify a unique edge $e_j \in p_n$, that contains the point. Furthermore, we can obtain the local y -coordinate as $y_0 = T - y_j^-$, where y_j^- is the y component of the coordinates of the lower vertex. We can find x_0 as $x_0 = \frac{y_0}{|e_j|} \Re(d_j)$.

- (2) An infinite number of intersection points; In this case, by Proposition 6.6 we have only at most $n - 1$ points (in local coordinates), one each in the intersecting edges of p_n . Let the edges that lie on the OL be $e_{j_1}, e_{j_2}, \dots, e_{j_h}$, where $h \leq n - 1$. In this case, we have

$$y_{j_1}^- = \dots = y_{j_h}^- = T,$$

so that $y_0 = 0$ in local coordinates for all these h edges. We also have $x_0 = 0$ for all these edges. Thus, in this case we have h root nodes.

REMARK 7.4. *In order to capture both these cases uniformly, one could create a dummy node e_0 and connect it to the edges of p_n that intersect the OL, by suitable arc weights. Thus, we can assume, without loss of generality, that there is a single node $e_0 \in p_n$ that contains the intersection point whose local coordinates are $z_0 = x_0 + \iota y_0$.*

7.2 The Graph G_0

We let $G_0 := (V_0, E_0)$ be the layered directed acyclic graph corresponding to Equation (7.1), where V_0 is the set of nodes, and E_0 the set of arcs. We denote the number of nodes by $|V_0|$, and the number of arcs by $|E_0|$. The layers of the graph correspond to levels $0, 1, 2, \dots, n$. The distinguished node $e_0 \in V_0$ is called the *root node* which is in level zero.

DEFINITION 7.5. *(Node in graph G_0) A node in the graph G_0 , is an edge $e_j \in p_k$ (or q_k) for some $k \in [n]$, with lower vertex z_j^- and higher vertex z_j^+ measured relative to the origin of subspace S_k containing p_k and q_k .*

DEFINITION 7.6. *(Arc in graph G_0) An arc in the graph G_0 connects a pair of edges e_i and e_j if and only if they satisfy the following conditions:*

- (i) $e_i \in p_k$ and $e_j \in q_k$, or vice versa, for some $k \in [n]$,
- (ii) they are on consecutive layers, i.e., $\text{level}(e_j) = \text{level}(e_i) \pm 1$, and
- (iii) and they interact, i.e., $|e_i \cap e_j| \geq 0$.

Furthermore, we associate the complex number $\overline{w}_{ij}$ to the arc, where $w_{ij} = z_j^- - z_i^-$, and call it the arc weight.

7.3 Computation of arcs between complementary curves

If we look at Figure 7 we see arcs going from p type curves to q type curves and vice versa alternatingly. Every arc in G_0 is an arc going from some $e_i \in p_k$ to some $e_j \in q_k$ or vice versa, where $k \in [n]$. Thus, if we compute all the arcs corresponding to interacting edges in all the n pairs of complementary curves, we can copy them into the various layers, and populate the graph. Let edge $e_i \in p_k$ and $e_j \in q_k$, and if $|e_i \cap e_j| \geq 0$, then the arc weight from e_i to e_j is simply $w_{ij} = z_j^- - z_i^-$. Similarly, the arc weight from e_j to e_i is simply the negation, i.e., $w_{ji} = -w_{ij}$. For the purpose of copying, it is useful to represent the arcs by an *adjacency list*, defined below.

DEFINITION 7.7. (*Adjacency List for an edge $e_i \in p_k$*) Let $k \in [n]$ and $i \in [\frac{k(k+1)}{2}]$. To each edge $e_i \in p_k$, we associate a list of pairs, denoted by $\mathcal{A}(p, k, i)$ and defined as

$$\mathcal{A}(p, k, i) := \{(r_1, w_{ir_1}), (r_2, w_{ir_2}), \dots\},$$

called the adjacency list, where the first element of each pair represents the edge in q_k that interacts with e_i , and the second element represents the complex arc weight. Similarly for an edge $e_i \in q_k$, the adjacency list is denoted by $\mathcal{A}(q, k, i)$.

By checking for interactions between all edges of a pair of complementary curves p_k and q_k , and for all $k \in [n]$, we can compute all the $|P_n| + |Q_n|$ adjacency lists. The complexity of this operation is at most $O(n^4)$, since there are n pairs of complementary curves, and for each pair we need to check at most $\frac{n^3+6n^2-n}{6}$ interactions by Proposition 6.13.

7.4 Taking character sets into account: appending adjacency lists

For an edge $e_i \in q_k$ (or p_k), we have seen that it can have at most k local characters (including itself). For example, an edge $e_i \in q_k$, can be congruent to edges $e_{i_1} \in q_1, \dots, e_{i_{k-1}} \in q_{k-1}$. In our compact graph Equation (7.1), since we collapsed all the λ transforms and have only τ transformations, we need to append the adjacency lists of the congruent edges on lower indexed curves to the adjacency list of e_i .

DEFINITION 7.8. (*Collated adjacency lists*) Let $e_i \in p_k$ for some $k \in [n]$ with adjacency list $\mathcal{A}(p, k, i)$. Let the character set of e_i be $\Lambda(e_i) = [r]$ where $r \leq k$. This means that the edges $e_{i_1} \in p_1, \dots, e_{i_r} \in p_r$ are congruent to e_i by λ transforms. We append the adjacency lists of these edges to that of e_i to get

$$\bar{\mathcal{A}}(p, k, i) := \bigcup_{\substack{e_{i_u} \in p_u \\ u \in [r], e_{i_u} \cong e_i}} \mathcal{A}(p, u, i_u)$$

to get a collated adjacency list $\bar{\mathcal{A}}(p, k, i)$.

After computing the initial adjacency lists for all edges in $\{P_n, Q_n\}$, we obtain the collated adjacency lists by taking into account the local characters of each edge. The complexity of this operation is simply $O(|P_n| + |Q_n|) = O(n^3)$.

7.5 Computation of nodes

The nodes can be computed easily since they are simply edges of the curves $\{P_n, Q_n\}$. For each edge $e_i \in p_k$, where $k \in [n]$ and $i \in [k(k+1)/2]$, recall that

$$z_i^- = \phi_k(i) + \iota \sigma_{\mathbf{a}}(\phi_k(i)),$$

$$\text{and } z_i^+ = \phi_k(i+1) + \iota \sigma_{\mathbf{a}}(\phi_k(i+1)).$$

Similarly for edges on q_k curve which make use of φ_k sequence. We copy these edges in layers as in Equation (7.1) and assign an unique index to each node. Then for each

node in G_0 , we connect to nodes in the next layer using the arcs given by the collated adjacency lists $\bar{\mathcal{A}}$. This completes the construction of graph G_0 .

Using the above description, we get a directed acyclic graph $G_0 = (V_0, E_0)$, where $V_0 = V(G_0)$ is the set of nodes, and $E_0 = E(G_0)$ is the set of arcs. We have designated root node $e_0 \in V_0$, and a start point $z_0 \in e_0$ given in local coordinates.

7.6 Size of G_0

The total number of nodes in the graph G_0 , (including the root node $e_0 \in p_{n+1}$) can be calculated (see Equation (7.1)) as

$$\begin{aligned} |V_0| &= 1 + \sum_{j=1}^n |Q_j| \\ &= 1 + \sum_{j=1}^n \frac{j(j+1)(j+2)}{6} \text{ (by Proposition 7.2)} \\ &= 1 + \frac{n(n+1)(n+2)(n+3)}{24} \approx \frac{n^4}{24}, \text{ for large enough } n. \end{aligned}$$

To compute the arcs consider one layer $Q_i \xrightarrow{\tau} P_{i-1}$ first. We know from Proposition 7.3 that

$$\#arcs \text{ in } Q_i \xrightarrow{\tau} P_{i-1} \approx \frac{i^5}{120}.$$

Now summing over all the layers, and adding the first layer arcs between the root node in p_n and Q_n (the second term in the expression below), we have

$$\begin{aligned} |E_0| &\leq \sum_{i=3}^n \frac{i^5}{120} \text{ (by Proposition 7.3)} \\ &\quad + \frac{n^3}{6} \text{ (by connecting root node to all nodes of first layer)} \\ &\approx \frac{n^6}{720}, \text{ for } n \text{ large enough.} \end{aligned}$$

For a random instance, if we use the estimate of Proposition 6.13 (which is more realistic), then the total number of arcs will be $\approx n^5/60$, which is a factor of $n/12$ less than the above estimate.

For later reference, we have the definition of a *leaf node* given below.

DEFINITION 7.9. (*Leaf node*) A node $e_k \in G_0$, is said to be a leaf node if and only if there are no outgoing arcs from e_k into the next level of the graph. Since each level of the graph corresponds to the family of curves Q_j or P_j for some $j \in [n]$, and recalling that nodes in p_4 and q_4 curves cannot have any children since 4 is the smallest local character allowed, it follows that leaf nodes can exist in all levels of the graph G_0 .

7.7 Paths in G_0

Given an instance $(\mathbf{a}, T)$, the previous section showed in detail the construction of a directed acyclic graph $G_0 = (V_0, E_0)$, with a start node e_0 , and a start value $z_0 = x_0 + \iota y_0 \in e_0$.

In this section we discuss paths in G_0 , their geometrical meaning and their classification.

DEFINITION 7.10. (*Path in G_0*) A path in G_0 is an ordered sequence of distinct nodes $(e_{k_1}, e_{k_2}, \dots, e_{k_l})$, such that every pair of adjacent nodes e_{k_i} and $e_{k_{i+1}}$ are on adjacent levels, and connected by a directed arc with weight $w_{k_i k_{i+1}}$, for all $1 \leq i < l$. We denote the path by $\pi(e_{k_1}, e_{k_l})$, identifying it by the first and last nodes. Sometimes we simply refer to a path by π when the start and end nodes are clear. The path length of the path $\pi(e_{k_1}, e_{k_l})$, denoted by $|\pi(e_{k_1}, e_{k_l})|$ is given by

$$|\pi(e_{k_1}, e_{k_l})| = \sum_{i=1}^{l-1} \bar{w}_{k_i k_{i+1}}.$$

DEFINITION 7.11. (*Image point*) Let $\pi(e_{k_1}, e_{k_l})$ be a path from e_{k_1} to e_{k_l} . Given a point $z_1 \in e_{k_1}$, the image of the point in any edge e_{k_j} , where $j \in [l]$, along the path is obtained by traversing the path up to the node, and is given by

$$z_j = z_1 + |\pi(e_{k_1}, e_{k_j})| = z_1 + \sum_{i=1}^{j-1} \bar{w}_{k_i k_{i+1}}.$$

Note that if $z_1 = 0$, then the image point is simply the path length.

DEFINITION 7.12. (*Valid path through e_i*) Let $\pi(e_0, e_i)$ be a path in G_0 , with start point $z_0 \in e_0$. The image point in e_i is $z_i = z_0 + |\pi(e_0, e_i)|$. We say π is a valid path into e_i , if and only if $y_i \in [0, |e_i|)$. Otherwise, it is called an invalid path.

DEFINITION 7.13. (*Valid path*) The path π is said to be a valid path if and only if it is valid at each of the nodes along the path.

The geometrical meaning of a path π in G_0 is fairly intuitive and as follows: Consider a path into a node e_i from $z_0 \in e_0$. If the path length $y_i < 0$, then geometrically this means that e_i does not intersect the OL, and the lower vertex of e_i is *above* the OL. Similarly, if $y_i > |e_i|$, then e_i as a whole, is *below* the OL. If $y_i \in [0, |e_i|)$, then the OL intersects this instance of edge e_i .

DEFINITION 7.14. (*Zero Path*) Let $\pi(e_0, e_k)$ be a path from e_0 to e_k , and let e_i be an intermediate node along the path. For a given point $z_0 \in e_0$, we say that the path from e_0 to e_i , i.e., $\pi(e_0, e_i)$ is a zero path if and only if it is a valid path, and $y_i = 0$ or $y_j = 0$ for some $j \in [i+1, k]$. In the former case, the OL passes through the lower vertex of e_i . In the latter case, although the path may not pass through the lower vertex of e_i , it is still called a zero path, since it leads to a zero path length further along the path.

DEFINITION 7.15. (*Distinct Paths through e_i*) Let $e_i \in V_0 \setminus \{e_0\}$ be an arbitrary node. Let π and π' be two paths through e_i coming from the start point $z_0 \in e_0$. Let the image points in e_i corresponding to the two paths be y_i and y'_i respectively. We say that the two paths π and π' are distinct if and only if $y_i \neq y'_i$.

DEFINITION 7.16. (*Distinct Valid Paths and Distinct Zero Paths*) Let $e_i \in V_0 \setminus \{e_0\}$ be an arbitrary node. If π and π' are two distinct paths through e_i , and furthermore they both are valid, then we call them distinct valid paths. Similarly, if π and π' are two distinct paths through e_i , and furthermore they both are zero paths, then we call them distinct zero paths.

7.8 Classification of paths through a node

Consider G_0 with $z_0 \in e_0$ as the starting point, and let e_i be an arbitrary node at level $l > 0$ from e_0 . If we compute all paths in G_0 , in general, we will have many paths through e_i . We have the following quantities of interest:

- (i) $DZP_i :=$ the set of all distinct zero paths through e_i . There could be many paths through e_i that go to a destination node (including e_i itself) and meet it at the vertex point. Out of these we only consider those which have a unique intersection in e_i . In other words, if more than one zero path goes through e_i at the same point in e_i , we will consider it as one. The cardinality of this set is denoted by $|DZP_i|$.
- (ii) $ZP_i :=$ the set of all zero paths through e_i . Same as DZP_i but with multiplicities counted. It is clear that

$$|DZP_i| \leq |ZP_i|. \quad (7.2)$$

- (iii) $DVP_i :=$ the set of all distinct valid paths through e_i , i.e., paths in $[0, |e_i|)$, and are distinct. These include both zero and non-zero paths, and thus

$$|DZP_i| \leq |DVP_i|. \quad (7.3)$$

If $|e_i| > 0$ then it is clear that $|DVP_i| \leq |e_i|$. In the special case of $|e_i| = 0$, we have $|DVP_i| \leq 1$ since there could be a path through the lower vertex.

With the above definitions of paths in G_0 , the *decision version* of *Subset Sum Problem* reduces to deciding the existence of zero paths in G_0 . Formally, we have the following theorem.

THEOREM 7.17. Let $(\mathbf{a}, T)$ be a given instance of SSP and let the OL intersect some node $e_0 \in p_n$ with local coordinates (x_0, y_0) . Then the given instance has a solution if and only if there is at least one zero path π from e_0 to some $e_k \in V_0$.

8 Hopping On The Orbital Line: An Algorithm

In this section we present an iterative algorithm to determine the number of *zero paths*, or alternatively the number of *vertex points* on the OL. The algorithm consists of repeatedly applying the alternating steps of *refining* and *filtering* the graph, which are described below. At the end, we have a refined graph G_m , which if not empty, has the property that *every* path is a *zero path*. In the case where there is no solution to the given instance of the *Subset Sum Problem*, the final graph will be empty, i.e., $G_m = \emptyset$.

First, let's consider briefly on how one might find zero paths. A naive approach would be to compute all paths from $z_0 \in e_0$ to all other nodes. However, this would be wasteful since zero paths are expected to be a small subset, if any at all, of the set of all valid paths. Furthermore, it is not clear what the complexity of finding all valid paths would be (we will later reveal this complexity too in next section). Instead, we do this: in each iteration of the algorithm, we make use of a routine called *Single Source Shortest Paths* abbreviated as SSSP, to find the shortest paths to all nodes in the graph. While it is not the same as finding zero paths, when combined with the refining and filtering steps, it suffices. So we first review this well-known graph algorithm [8; 21] which has a linear running time in the size of the graph.

8.1 Single Source Shortest and Longest Paths

It is a well-known fact that every directed acyclic graph admits a *topological ordering* of its nodes. Given a directed acyclic graph $G = (V, E)$, the topological ordering of the graph is simply a linear ordering of the nodes on a horizontal line so that all the arcs go from left to right. Furthermore, the topological ordering can be accomplished in linear time in the size of the graph, i.e., in time $O(|V| + |E|)$, using a *depth first search* [8].

The SSSP algorithm is slightly modified so that we also have the *longest* path to each node. The algorithm for computing the *shortest* and the *longest* paths from the root node to all other nodes, in a directed acyclic graph, is given below (Algorithm 1). Here, we assume that we have already carried out the topological ordering and that the ordered nodes are stored in an array labeled *TS*.

Given that the arc weights (in the y component) in the graph can be negative, in general, this algorithm may output negative shortest paths, and large positive paths (greater than the edge length) to many nodes.

NOTATION 8.1. For any node $e_i \in V_0$, we denote the shortest path from e_0 by α_i^* and the longest path by $\alpha_i^\square$, where $\alpha_i^*, \alpha_i^\square \in \mathbb{Z}$.

8.2 Refining Nodes and Arcs

In this subsection, we define the notion of *refining* of nodes and arcs, and the generation of a refined graph.

DEFINITION 8.2. (*Refine operation on a node*) Let e_i be a node with the associated interval $[z_i^-, z_i^+)$, and with length $|e_i| = y_i^+ - y_i^- > 1$ where $i \in [|V_0|]$. Let $z_i^h = \lfloor \frac{(z_i^- + z_i^+)}{2} \rfloor$

Algorithm 1 Pseudo-code to compute the shortest and longest paths to all nodes from root, given the topologically sorted order of nodes (TS), and initial value $z_0 \in e_0$ in the root node.

```

1: procedure SSSP( $G = (V, E)$ ,  $TS$ ,  $e_0$ ,  $z_0$ ,  $\alpha^*$ ,  $\alpha^\square$ )
2:    $\alpha_i^* \leftarrow \infty$  for all  $i \in [|V|]$  ▷ initialize all shortest paths to  $\infty$ 
3:    $\alpha_i^\square \leftarrow -\infty$  for all  $i \in [|V|]$  ▷ initialize all longest paths to  $-\infty$ 
4:    $\alpha_0^* \leftarrow y_0$  ▷ set root shortest path to  $y_0$ 
5:    $\alpha_0^\square \leftarrow y_0$  ▷ set root longest path to  $y_0$ 
6:   for  $j = 2$  to  $|V|$  do
7:      $u \leftarrow TS[j]$  ▷ for each node in the topological order
8:     for  $k \in P(u)$  do ▷ For each parent of  $u$ 
9:       if  $\alpha_k^* + \bar{w}_{ku} < \alpha_u^*$  then
10:         $\alpha_u^* \leftarrow \alpha_k^* + \bar{w}_{ku}$ 
11:       if  $\alpha_k^\square + \bar{w}_{ku} > \alpha_u^\square$  then
12:         $\alpha_u^\square \leftarrow \alpha_k^\square + \bar{w}_{ku}$ 
13:   return  $\alpha^*, \alpha^\square$  ▷ Return the shortest and longest paths

```

be the midpoint of the interval. A refine operation on e_i , replaces it with two smaller disjoint nodes e'_i and e''_i whose union is e_i . In particular, e'_i corresponds to the interval $[z_i^-, z_i^h)$, and e''_i corresponds to $[z_i^h, z_i^+)$.

REMARK 8.3. For a node e_i , if the length $|e_i| \leq 1$, we do not split it into two nodes. We simply leave it as it is. The floor operation ensures that the end points of nodes (in y coordinate) before and after refining remain as integers.

Consider a pair of nodes e_i and e_j connected by an arc of weight w_{ij} . Recall that the arc weight is determined by the lower vertices of the two nodes, i.e., $w_{ij} = z_j^- - z_i^-$. When we *simultaneously* refine the nodes e_i and e_j , we have four new nodes, and have potentially four possible arcs between them. See Figure 12 below for an illustration.

The refine operation on e_i leads to new edges e'_i and e''_i . Similarly, e'_j and e''_j are generated from e_j . The end points associated with these new edges are as follows:

$$\begin{aligned}
e'_i &\Leftrightarrow [z_i^-, \lfloor \frac{z_i^- + z_i^+}{2} \rfloor], \\
e''_i &\Leftrightarrow [\lfloor \frac{z_i^- + z_i^+}{2} \rfloor, z_i^+), \\
e'_j &\Leftrightarrow [z_j^-, \lfloor \frac{z_j^- + z_j^+}{2} \rfloor], \\
\text{and } e''_j &\Leftrightarrow [\lfloor \frac{z_j^- + z_j^+}{2} \rfloor, z_j^+).
\end{aligned}$$




Figure 12: Illustration of refining edges e_i and e_j connected by an arc. Only three pairs of new edges interact; edges e_j' and e_i'' do not interact since their y intervals do not overlap.

We have four possible arcs between the new nodes, with arc weights as follows:

$$\begin{aligned}
 (e_i', e_j') &\Leftrightarrow w_{ij}, \\
 (e_i', e_j'') &\Leftrightarrow w_{ij} + \lfloor \frac{z_j^+ - z_j^-}{2} \rfloor, \\
 (e_i'', e_j') &\Leftrightarrow w_{ij} - \lfloor \frac{z_i^+ - z_i^-}{2} \rfloor, \\
 \text{and } (e_i'', e_j'') &\Leftrightarrow w_{ij} + \lfloor \frac{z_j^+ - z_j^-}{2} \rfloor - \lfloor \frac{z_i^+ - z_i^-}{2} \rfloor.
 \end{aligned}$$

However, the proposition below shows that we will have at most three interacting pairs of edges and never four.

PROPOSITION 8.4. *Let e_i and e_j be two nodes with length greater than unity each, and connected by an arc with weight w_{ij} . A simultaneous refine operation on e_i and e_j will generate at most three pairs of interacting edges between $\{e_i', e_i''\}$ and $\{e_j', e_j''\}$.*

We omit the proof as it is a simple exercise to prove it by considering various cases.

COROLLARY 8.5. *Let $G_0 = (V_0, E_0)$ be the initial graph. Then after simultaneous refinement of all nodes, and connecting interacting pairs of nodes by arcs, the refined graph $G_1 = (V_1, E_1)$ satisfies:*

- (i) *the node with maximum length, say e_i , has $|e_i| \leq 2^{m-1}$,*
- (ii) *$|V_1| \leq 2|V_0|$,*
- (iii) *and $|E_1| \leq 3|E_0|$.*

The complexity of the *refine* algorithm is clearly $O(|V| + |E|)$ since it involves one loop over nodes and one loop over arcs. The check for interactions can be done in

constant time. Since the source node e_0 is also split into two halves, the original point $y_0 \in e_0$ can be in either e'_0 or e''_0 but not both. In case it is in e'_0 , the (local) value of y_0 does not change. Otherwise, we update the local value in the reference frame of e''_0 as $y_0 := y_0 - \frac{|e_0|}{2}$. The half containing the start point is kept and relabeled e_0 , and the other half is removed. We can find all the nodes in the refined graph that are *reachable* from e_0 by doing a simple *Depth First Search (DFS)* [8]. This step will get rid of some unreachable nodes and arcs of the refined graph.

The pseudo-code for the *refine* operation is shown below (Algorithm 2). The complexity is clearly linear in the size of the input graph, namely $O(|V| + |E|)$.

Algorithm 2 (Pseudo-code for the Refine operation)

```

1: procedure REFINE( $G = (V, E)$ ,  $e_0$ ,  $z_0$ ,  $G'$ )
2:    $G' \leftarrow \emptyset$  ▷ Initial graph  $G'$  to  $\emptyset$ 
3:   for  $i = 1$  to  $|V|$  do
4:      $e_i \rightarrow \{e'_i, e''_i\}$  ▷ Apply the refine operation
5:     Add  $\{e'_i, e''_i\}$  to  $V'$ 
6:     for  $(i, j) \in E$  do
7:       for  $r \in \{', ''\}$  do
8:         for  $s \in \{', ''\}$  do
9:           if  $e_i^r \cap e_j^s \neq \emptyset$  then ▷ Check if edges interact
10:            Add arc  $(e_i^r, e_j^s)$  to  $E'$ 
11:   if  $z_0 \in e'_0$  then
12:      $e_0 \leftarrow e'_0$ 
13:   else
14:      $e_0 \leftarrow e''_0$  ▷ update new root node
15:      $z_0 \leftarrow z_0 - \lfloor \frac{z_0^- + z_0^+}{2} \rfloor$  ▷ Update value of  $z_0$  in new reference frame
16:    $G' = REACH(G', e_0)$  ▷ Get Reachable DAG from  $e_0$  by DFS
17:   return  $G'$  ▷ Return the refined graph  $G'$ 

```

8.3 A destination node e_∞

In the graph G_0 we have a distinguished *root node* labeled e_0 . In order to facilitate the *filtering* operation it is useful to have a destination node e_∞ defined below.

DEFINITION 8.6. (*TRUE and FALSE nodes*) Let $e_i \in G_0$ be a node, with end points z_i^- and z_i^+ . The node e_i is said to be a *TRUE* node if and only if the lower vertex belongs to the point set, i.e., $z_i^- \in S_n$. Otherwise it is called a *FALSE* node. In the initial graph G_0 , all nodes are *TRUE* nodes. However, after applying a refine operation, only the lower-half nodes will be *TRUE* nodes. The upper-half nodes will be *FALSE* since their lower vertices are midpoints of the segments which are not part of the point set.

DEFINITION 8.7. (*destination node*) We include a dummy node labeled as e_∞ , and attach it to all TRUE nodes of $V_0 \setminus \{e_0\}$ by arcs with arc weight equal to zero. As a result of this operation, the number of arcs in the modified graph is equal to $|E_0| + |V_0| - 1$. Note that while the rest of the nodes in the graph have a geometrical interpretation, the dummy node e_∞ does not have any.

The following proposition is straightforward, given the above definitions.

PROPOSITION 8.8. Given G_0 , and a $z_0 \in e_0$, there is a zero path to some node $e_k \in V_0$, if and only if there is a zero path to e_∞ .

COROLLARY 8.9. Given G_0 , and a $z_0 \in e_0$, if there are N zero paths to various nodes, then there will be N zero paths to e_∞ .

In a sense, the node e_∞ acts like a collector of all zero paths.

DEFINITION 8.10. (*Reverse graph $\hat{G}_0$*) Let $\hat{G}_0$ be the graph obtained by reversing all arcs, with e_∞ as the root node, and e_0 as the destination node. Note that each arc w_{ij} of the original graph, becomes $-w_{ij}$ in the reverse graph. We can start with $z_\infty = 0$ as the start point in e_∞ , and propagate paths back to e_0 . Let β_k^* and $\beta_k^\square$ denote the shortest and longest paths to a node e_k in the reverse graph.

8.4 Filtering of nodes and arcs

The main idea in *filtering* is to remove nodes and arcs of the graph, that *cannot support* any zero paths. In other words, if no zero paths can go through a node (or an arc), they can be safely removed. We have already seen above that, if all nodes of G_0 are refined, then for every pair of nodes connected by an arc in the original graph, we will have only at most three arcs instead of four in the refined graph. In a sense, the refine operation also filters some arcs. If we view the nodes connected by arcs as *channels*, the width is reduced, and they become more *aligned*. Here, we describe some tests which will enable us to remove more nodes and arcs.

Given $z_0 \in e_0$, first we compute the *shortest* and *longest* paths to all other nodes in the graph using the SSSP algorithm. Note that we are only concerned with the y coordinate in the computation of paths. The x coordinate gets a free ride!

PROPOSITION 8.11. Let α_k be an arbitrary path into a node $e_k \in G_0$. Then, we have

$$\alpha_k^* \leq \alpha_k \leq \alpha_k^\square.$$

This simply follows from the fact that α_k^* and $\alpha_k^\square$ are the extreme shortest and longest paths respectively into e_k .

Definition 8.6 of TRUE/FALSE labels of nodes has an important use. A TRUE node has a vertex point, and thus it has the *potential* to receive a zero path, whereas a FALSE node has none. Let e_k be a TRUE node. If it turns out that the shortest path $\alpha_k^* > 0$ then it can be seen that node e_k can never have a zero path into it. We can then relabel the node e_k as FALSE.

PROPOSITION 8.12. *Let $G = (V, E)$ be the OL graph at some stage of iteration, and let $e_k \in V \setminus \{e_0\}$ be a node in G . A necessary condition that at least one instance of e_k intersects the orbital line is*

$$[0, |e_k|) \cap [\alpha_k^*, \alpha_k^\square] \neq \emptyset$$

Proof. Recall that each path into a node e_k corresponds to a sequence of line segments in the geometry. This sequence of segments has e_0 as the first segment and e_k as the final segment. In order for this instance of e_k to intersect the OL, the path length must be in $[0, |e_k|)$. Since every path into e_k lies in the interval formed by the shortest and longest paths into e_k , the condition is clear. Suppose all instances of e_k are above the OL. Then the shortest and longest paths are both negative and the above condition leads to an empty interval. Similarly, if all instances of e_k are below the OL, then the shortest and longest paths will be greater than the edge length. Again, the above condition leads to an empty set. $\square$

Now we are ready to describe *Filtering* of the graph.

PROPOSITION 8.13. *(Node Removal) Let $G = (V, E)$ be the OL graph at some stage of iteration, and let $e_k \in V \setminus \{e_0, e_\infty\}$ be a node in G . Then e_k can be removed if either of the following conditions is satisfied.*

- (i) *The node e_k is not reachable from e_0*
- (ii) *The forward and reverse shortest and longest paths into e_k satisfy*

$$[0, |e_k|) \cap [\alpha_k^*, \alpha_k^\square] \cap [\beta_k^*, \beta_k^\square] = \emptyset$$

Proof. Condition (i) is clear: in order for e_k to be in the graph, it must be reachable from e_0 via edge interactions. Coming to the first two terms in LHS of condition (ii), we need at least one instance of e_k to intersect the OL, as a necessary condition, by Proposition 8.12. The third term on LHS of (ii) is similar to the middle term, but in the reverse direction. If there is a zero path passing through node e_k reaching e_∞ , then we need a collision of forward and reverse paths at e_k . A necessary condition for this is that these forward and reverse intervals have a non-empty intersection, and the claim follows. $\square$

PROPOSITION 8.14. *(Arc Removal) Let $G = (V, E)$ be the OL graph at some stage of iteration, and let $w_{ij} \in E$ be an arc. Then the arc w_{ij} can be removed if either of the following conditions is satisfied.*

- (i) *the intervals of shortest and longest paths, in the forward direction, of e_i and e_j do not interact*

$$[\alpha_i^* + \overline{w}_{ij}, \alpha_i^\square + \overline{w}_{ij}] \cap [\alpha_j^*, \alpha_j^\square] = \emptyset$$




Figure 13: Collision of forward and reverse paths at e_k

- (ii) the intervals of shortest and longest paths, in the reverse direction, of e_i and e_j do not interact

$$[\beta_j^* - \overline{w}_{ij}, \beta_j^\square - \overline{w}_{ij}] \cap [\beta_i^*, \beta_i^\square] = \emptyset$$

Proof. The basic idea is this: although nodes e_i and e_j interact (with their edge intervals), they may not interact when we use the intervals formed by shortest and longest paths. In such a case, it is straightforward to see that the shortest path at node e_i cannot lead to shortest path at e_j . Hence, we can safely remove this arc. $\square$

Algorithm 3 (Pseudo-code for Filter operation)

```

1: procedure FILTER( $G = (V, E)$ ,  $e_0$ ,  $z_0$ ,  $G'$ )
2:    $G' \leftarrow G$  ▷ Copy the initial graph
3:   for  $i = 1$  to  $n$  do ▷ run this loop at most  $n$  times
4:     Attach destination node  $e_\infty$ 
5:      $G' = \text{REACH}(G', e_0, e_\infty)$  ▷ reachable graph between  $e_0$  and  $e_\infty$ 
6:      $z_\infty \leftarrow 0$ 
7:     Compute  $\alpha^*, \alpha^\square, \beta^*$  and  $\beta^\square$  ▷ Compute extreme paths
8:     for  $j = 1$  to  $|V|$  do
9:       if  $[0, |e_j|] \cap [\alpha_j^*, \alpha_j^\square] \cap [\beta_j^*, \beta_j^\square] = \emptyset$  then
10:        Mark  $e_j$  for deletion
11:       if  $\alpha_j^* > 0$  &  $e_j$  is a TRUE node then
12:        Set  $e_j$  as FALSE
13:       for  $k = 1$  to  $|E|$  do
14:         if  $[\alpha_i^* + \overline{w}_{ij}, \alpha_i^\square + \overline{w}_{ij}] \cap [\alpha_j^*, \alpha_j^\square] = \emptyset$  then
15:           Mark  $w_{ij}$  for deletion
16:         if  $[\beta_i^* - \overline{w}_{ij}, \beta_i^\square - \overline{w}_{ij}] \cap [\beta_j^*, \beta_j^\square] = \emptyset$  then
17:           Mark  $w_{ij}$  for deletion
18:       if (no marked nodes and arcs) then
19:         delete dummy node and break
20:       Delete marked nodes and arcs, including dummy node
21:   return  $G'$  ▷ Return the filtered graph  $G'$ 

```

The pseudocode for the filtering operation is shown in Algorithm 3. In the inner loop of filter operation, there are two calls to SSSP to compute the shortest and longest paths in both directions. To figure out the nodes and arcs for deletion, we loop over the nodes once, and over the arcs once. Again the overall complexity is linear in the size of the input graph. The loop is executed at most n times. The reason for this is: when some nodes and arcs are removed, the shortest and longest path lengths to the next level change. This is due to the greedy nature of the SSSP routine. After the k th execution of the loop, it can be shown that the shortest and longest paths at levels 1 through k will not change, so that after at most n iterations, we will not get any nodes and arcs that can be deleted.

REMARK 8.15. *At the end of each loop of the FILTER step, each surviving node e_i satisfies one of the following properties.*

- (i) *The node e_i has a zero path into it that has been discovered ($\alpha_i^* = 0$).*
- (ii) *The node e_i has a potential for a zero path into it ($\alpha_i^* < 0$ and $\alpha_i^\square > 0$) but it is not clear now. This uncertainty will be removed after further filter or refine steps.*
- (iii) *The node e_i does not have a zero path into it ($\alpha_i^* > 0$), but it can potentially support a zero path at a downstream node. In this case e_i is labeled as FALSE. If furthermore, e_i is a LEAF node, it will be disconnected from the destination node e_∞ in the next loop and removed (see step 5 of Algorithm 3). Otherwise, it will remain for now.*
- (iv) *When a node e_i that is both a LEAF node and a FALSE node is removed, one or more of its immediate parent nodes can become a LEAF node. If such a node is also FALSE, it will be removed in the next loop by above step (iii). Thus, there can be a cascaded removal of nodes (during the n loops of the FILTER operation) that do not support zero paths.*

As the iterations of the algorithm progress, the number of TRUE nodes is non-increasing. Meanwhile the FALSE nodes which do not support zero paths get removed.

8.5 The algorithm at a glance

Now we are ready to describe the algorithm. With the above detailed description of the *refine* and *filter* operations, the algorithm (referred to as IHM) is quite simple as shown in Algorithm 4. The algorithm consists of m iterations, where in each iteration we perform the *refine* and *filter* operations.

We have the progression of graphs

$$G_0, G_1, G_2, \dots, G_m,$$

as we go through the iterations. The refine operation doubles the nodes in each step, whereas the filter operation removes some nodes in each step. If the filter step weren't

Algorithm 4 High level description of the algorithm.

```

1: procedure IHM( $G_0 = (V_0, E_0)$ ,  $e_0$ ,  $z_0$ ,  $m$ ,  $G_m$ )
2:    $G \leftarrow G_0$  ▷ Copy the initial graph
3:    $G \leftarrow \text{FILTER}(G, e_0, z_0)$  ▷ Apply the filter operation
4:   for  $j = 1$  to  $m$  do
5:      $G' \leftarrow \text{REFINE}(G, e_0, z_0)$  ▷ Apply the refine operation
6:      $G_j \leftarrow \text{FILTER}(G', e_0, z_0)$  ▷ Apply the filter operation
7:      $G \leftarrow G_j$ 
8:     if  $G_j = \emptyset$  then
9:       return  $\emptyset$ 
10:  return  $G_m$  ▷ Return the final graph  $G_m$ 

```

there we would trivially have

$$|V_m| \leq 2^m |V_0|,$$

$$\text{and } |E_m| \leq 3^m |E_0|,$$

making the final graph exponential in m .

How many nodes does the filter operation remove? It can be shown quite easily that the filter operation removes at least a polynomial number (in n) of nodes, simply by noting that there are $O(n^3)$ nodes in the first level of the orbital graph, but only $O(n)$ valid paths into them. To properly quantify the growth of the graph, we have to wait for the next section.

8.6 Growth factor

Meanwhile, we define a simple metric to measure the growth rate. For this, we simply use the number of nodes as a measure. Let

$$\eta_i := \frac{|V_i|}{|V_0|}, \text{ for } i \in [m],$$

be the growth factor as a function of iterations.

DEFINITION 8.16. (*peak growth factor*) Let η_{peak} denote the peak growth factor of nodes in IHM. Then

$$\eta_{peak} := \max_{i \in [m]} \eta_i = \frac{\max_{i \in [m]} |V_i|}{|V_0|}.$$

PROPOSITION 8.17. The complexity of Algorithm 4 is $O(\eta_{peak} m n^7)$.

Proof. First notice that at any iteration of the algorithm, the graph size is $O(\eta_{peak} |G_0|) = O(\eta_{peak} n^6)$. The refine operation is simply linear in the size of the graph. The filter operation is also linear in the size of the graph, but executed at most n times. Thus, the complexity per iteration is $O(\eta_{peak} n^7)$. Since the algorithm consists of m iterations, the overall complexity is $O(\eta_{peak} m n^7)$. $\square$

We will obtain a bound for η_{peak} in the next section. It can be readily seen that η_{peak} depends on the *maximum number of distinct zero paths* ($\max DZP$). Regarding the nature of the final graph G_m , we have this:

THEOREM 8.18. *The graph $G_m = (V_m, E_m)$ has the following properties:*

- (i) *All nodes have length ≤ 1 .*
- (ii) *All arc lengths are zero in the y -component.*
- (iii) *All leaf nodes are TRUE nodes.*
- (iv) *All paths are zero paths.*

Proof. With each refine step, the node lengths are halved. Since the maximum possible initial node length is 2^m , in at most m iterations, all node lengths become either unity or zero. Since two nodes can be connected by an arc only if they interact, using the fact that all node lengths are unity or zero, it follows that arc weights in the y component are zero. This takes care of (i) and (ii).

Recall that a node $e_i \in V_m$ is a TRUE node, if $(x_i^-, y_i^-) \in S_n$. Also recall that a leaf node is a node with no outgoing arcs. If some leaf node e_i is a FALSE node, then Step 5 of the Algorithm 3, would remove it since it will not be connected to e_∞ . Thus, every leaf node in G_m is a TRUE node, and this takes care of (iii).

Finally, note that the initial point $y_0 = 0$ in the refined start node e_0 (since it is of unit length, and is a half-open interval). Since all arc weights are zero in the y component, it follows that all paths are zero paths, and (iv) is true. $\square$

DEFINITION 8.19. (*Solution Graph*) If $G_m \neq \emptyset$, we call it as the *Solution Graph*.

8.7 Indices as path lengths

Recall that in the local reference frame of an edge, the local x coordinate is determined from the local y coordinate (see Section 6.2). For a given OL, the intersection point in e_0 is $z_0 = x_0 + \iota y_0$, and $z_0 = z_0^- + x_0^{loc} + \iota y_0^{loc}$. In general, x_0^{loc} is a rational number. However, we can ignore this and set $x_0^{loc} = 0$, i.e., treat x_0 as x_0^- the global x coordinate of edge e_0 . We do this for every edge along the path. Since the x -component of the arc weight is simply the translation along x between the pair of edges, the overall the path length gives the x coordinate of the point in the geometry. With this provision, we will have only integer values for x and the following result.

THEOREM 8.20. *Let π be a zero path in G_m starting at e_0 and ending at some TRUE node e_k . Then the x -component of the path length, denoted $|\pi|_x$, equals the index of the subset whose subset sum is the target T .*

9 A Holy Trinity: Number, Size and Additive Structure

The purpose of this section is to provide an upperbound on the maximum number of distinct zero paths through any equivalence class of edges intersecting the OL. In order to do that let's take a contemplative look on the nature of the SSP itself! Recall that the Subset Sum Problem is specified by two independent parameters: the number of elements n , and the maximum size of elements in bits m . However, there is also a third quantity involving the relationship between the elements which we simply call as the *degree of additive structure* (to be defined later).

It is the interplay of these three quantities, which we call a *trinity of factors*, that determine the ease or difficulty of finding the solution to a given instance of the SSP. In this section, we examine these three quantities and reveal the precise way in which they control the complexity. Finally we provide an upperbound on the maximum number of locally distinct paths through any edge equivalence class intersecting the OL.

The general theme is that the geometric description of the solution space of SSP has two aspects:

- (i) a *combinatorial aspect* depending only on n ,
- (ii) and a *relational aspect* depending on m and the *degree of additive structure*.

Many of the quantities that fall out of symmetries such as the number of paths to level r , i.e., $\beta(r)$, are of *combinatorial* nature, and these are exponential in number. Then there are other quantities like *distinct path lengths* which have a *relational aspect* to them, and they depend on m and the *degree of additive structure*.

To determine $\max DZP$, we first consider *all* the intersection points on the *orbital line* and bound the *maximum distinct valid paths* ($\max DVP$) instead; this automatically provides a bound on $\max DZP$ since the set of all *zero paths* is a subset of the set of all *valid paths*. The main reason we look at this larger set is to understand the structure of points around the OL, and as the reader will soon see, analyzing this larger set also makes things simpler mathematically. For this, we define a *configuration graph* that captures both the *combinatorial* and the *relational* aspects of the problem. This is followed by a rigorous analysis of the number of paths in the configuration graph, from which an upperbound on the $\max DVP$ can be extracted. Then a bound for $\max DZP$ can be derived (which is a special case) easily.

As a consequence of this analysis, a curious finding that *size doesn't matter beyond a threshold value dependent on $\log n$* is obtained. We also provide a precise estimate for this critical size m_0 .

9.1 A Configuration Graph

Let $M := 2^m$. Consider a sequence of $n+1$ vertical segments denoted $(s_0, s_1, s_2, \dots, s_n)$, where

$$s_i := \{(x, y) \in \mathbb{R}^2 : x = i \in [0, n] \text{ and } y \in [0, M)\},$$

arranged along the x -axis and are parallel. Thus, each segment has 2^m integer points and these are of interest in what follows.

In the graph G_0 , given a start value $y_0 \in e_0$, let the set of all valid paths of length $r \in [n]$ on the *orbital line* be denoted by Π_r . For each path $\pi \in \Pi_r$, since it is valid, we have a sequence of edges $(e_{\pi_1}, e_{\pi_2}, \dots, e_{\pi_r})$ that it goes through with local intercepts $(y_{\pi_1}, \dots, y_{\pi_r})$ respectively, where $y_{\pi_i} \in [0, |e_{\pi_i}|]$ for all $i \in [r]$. We mark the point $(0, y_0)$ in the segment s_0 , $(1, y_{\pi_1})$ in segment s_1 , and so on with (r, y_{π_r}) marked in s_r . We connect points on adjacent segments by a straight line. With this we have a piece-wise linear curve corresponding to the path π . We call this the *configuration plot* of the path π . An example is shown below in Figure 14.




Figure 14: Configuration plot of a valid path π of length r along the OL. Although the edge lengths can vary (but are limited by $M = 2^m$), we use this maximum possible length to accomodate all possible local intercepts.

REMARK 9.1. For a path $\pi \in \Pi_r$, each segment in the configuration plot corresponds to some arc weight w_{ij} relating two interacting edges $e_i, e_j \in G_0$. While w_{ij} has both x and y components, we only care about the y component in the configuration plot, and it can be seen that it lies in the interval $[-M + 1, M - 1]$.

DEFINITION 9.2. (Locally Distinct Configurations) Two paths π and π' , both of length r , are said to have locally distinct configurations if and only if

$$(y_{\pi_1}, y_{\pi_2}, \dots, y_{\pi_r}) \neq (y'_{\pi_1}, y'_{\pi_2}, \dots, y'_{\pi_r}).$$

DEFINITION 9.3. (The set of locally distinct configurations Ξ_r) Let Ξ_r denote the set of all locally distinct configurations obtained from Π_r . In other words, this is a subset of Π_r that produce unique configurations. We denote the size of this set by $|\Xi_r|$.

Since we are dealing with local values in the configuration plot without regard to which edges they came from, two paths using distinct sequences of nodes in G_0 can still result in the same configuration. As a result, we have:

PROPOSITION 9.4. For any $r \in [n]$, we have $|\Xi_r| \leq |\Pi_r|$.

If we overlay the configuration plots of all the paths Π_r for all $r \in [n]$, in the same configuration graph, we will have many distinct piece-wise linear curves emanating from the point $(0, y_0) \in s_0$, going through the intermediate segments and arriving at some integer point in s_r . In general, this will be a dense thicket of paths corresponding to

paths on the OL. Each segment s_r for $r \in [n]$ captures all the local values of paths entering edges of level r in G_0 . The configuration plot thus provides a compact view of all the paths on OL.

DEFINITION 9.5. (*Point set F_r and γ_r*) For each $r \in [n]$, let F_r denote the set of distinct points in segment s_r that have at least one incoming path from segment s_{r-1} . Each point in F_r has coordinates of the form (r, y) where $y \in [0, 2^m)$. Let $\gamma_r := |F_r|$ denote the number of distinct points in segment r . We have the starting set F_0 which consists of a single point $(0, y_0)$, and thus $\gamma_0 = 1$. It is also clear that $\gamma_r \leq 2^m$, for all $r \in [n]$.

DEFINITION 9.6. (*Arc set $C(r, r+1)$*) For a non-negative integer $r \in [0, n-1]$ let $C(r, r+1)$ denote the set of distinct arcs between levels r and $r+1$. Each arc of $C(r, r+1)$ is a line segment connecting some point $(r, y) \in F_r$ with some other point $(r+1, y') \in F_{r+1}$. It is clear that

$$C(r, r+1) \subseteq F_r \times F_{r+1}.$$

We denote the size of $C(r, r+1)$ by $|C(r, r+1)|$.

DEFINITION 9.7. (*Layered Configuration Graph $\mathcal{C}$*) We view the points in F_r for $r \in [n]$ as nodes, and line segments $C(r, r+1)$ as arcs of the configuration graph for $r \in [0, n-1]$. We denote this configuration graph by $\mathcal{C}$. The single point in F_0 serves as a root node. Thus, $\mathcal{C}$ is a layered graph with $n+1$ levels and comprised of n bipartite layers. We denote the set of all nodes of $\mathcal{C}$ by $V(\mathcal{C})$ and the set of all arcs by $E(\mathcal{C})$ using standard notation for graphs.

From the above definitions, it follows that

$$\begin{aligned} \# \text{ nodes} &= |V(\mathcal{C})| = \gamma_0 + \gamma_1 + \gamma_2 + \dots + \gamma_n, \\ \text{and } \# \text{ arcs} &= |E(\mathcal{C})| = |C(0, 1)| + |C(1, 2)| + \dots + |C(n-1, n)|. \end{aligned}$$

DEFINITION 9.8. (*point number sequence γ*) We define $\gamma := (\gamma_1, \gamma_2, \dots, \gamma_n)$ as the sequence of number of distinct points in each segment and refer to it as the point number sequence.

QUESTION 9.9. We are interested in answering the following questions:

- (i) What is $\max(\gamma)$?
- (ii) What is the asymptotic profile of γ for large m ?

It can be seen easily that $\max DVP \leq \max(\gamma)$, and hence the importance of the first question.

9.2 Apparent paths in $\mathcal{C}$

Once all the valid paths on the OL are plotted in the configuration graph, we can view $\mathcal{C}$ as a graph on its own.

DEFINITION 9.10. (*Paths in $\mathcal{C}$*) A path of length $r \in [n]$ in $\mathcal{C}$ is a sequence $(0, y_0), (1, y_1), \dots, (r, y_r)$ of points one from each of the sets $F_0, F_1, \dots, F_r$, such that every two adjacent points are connected by an arc.

Note that a connected path in $\mathcal{C}$ may or may not correspond to a path in G_0 as the example plot below shows.




Figure 15: Illustration of real and apparent paths in $\mathcal{C}$.

In Figure 15, there are two paths $\pi_1 = (y_0, y_1, y_2, y_3, y_4)$ (shown in red) and $\pi_2 = (y_0, y'_1, y_2, y'_3, y_4)$ (shown in green) coming from the paths in Π_4 . However, two other paths are induced: $\pi_3 = (y_0, y_1, y_2, y'_3, y_4)$ and $\pi_4 = (y_0, y'_1, y_2, y_3, y_4)$. The paths π_3 and π_4 have no existence in G_0 , and hence are called *apparent paths*. Thus, the number of paths in $\mathcal{C}$ can be greater than what we started with!

9.3 Path magnification in $\mathcal{C}$

DEFINITION 9.11. (*out-degree*) For a given point $(r, y) \in F_r$, where $r \in [0, n-1]$, the number of distinct outgoing arcs from it to points in F_{r+1} is called the out-degree, and denoted by $\mu(r, y)$.

DEFINITION 9.12. (*Max out-degree*) Let μ_r denote the maximum out degree over all points in the r th level. Thus,

$$\mu_r := \max_{y \in F_r} \mu(r, y).$$

PROPOSITION 9.13. For a given level $r \in [0, n-1]$, the maximum out-degree μ_r is at most $\frac{(n-r)^3 + 3(n-r)^2 - 13(n-r) + 6}{3}$.

Proof. Suppose that level r corresponds to Q_{n+1-r} ; then level $r+1$ corresponds to P_{n-r} . The point $(r, y) \in F_r$ can belong to any of the $|Q_{n+1-r}|$ nodes of level r in G_0 . The transformation structure of paths going from level r to level $r+1$ is $q_j \xrightarrow{\lambda} q_k \xrightarrow{\rho} p_k$ where $j > k$. Each arc weight is between an edge of q_k and an edge of p_k as k ranges in $[n-r]$. Since we are interested in distinct arc weights, we consider the worst case where $\mathbf{a}$ is a

dissociated set. Since there are at most $k^2 + k - 5$ distinct arc weights between q_k and p_k (see Proposition 6.13), summing over $[n - r]$ we get

$$\begin{aligned}\mu_r &\leq \sum_{k=3}^{n-r} k^2 + k - 5 \\ &= \frac{(n-r)^3 + 3(n-r)^2 - 13(n-r) + 6}{3}\end{aligned}$$

and the claim follows. The result remains the same if level r corresponds to P_{n+1-r} , and level $r + 1$ corresponds to Q_{n-r} . $\square$

DEFINITION 9.14. Notice that μ_r is a decreasing function of r , and maximum value is attained at $r = 0$ where $\mu_0 = \frac{n^3+3n^2-13n+6}{3} \approx \frac{n^3}{3}$ for n large enough. We will be using $\mu_0 = \frac{n^3}{3}$ from here on.

REMARK 9.15. The above proposition regarding the out degree is essentially an estimate in the worst case. When the given instance has high degree of additive structure many arc weights will be the same. While these arcs will be distinct in G_0 , they will not be in $\mathcal{C}$ since we are only concerned with the y component of arc weight. As a result, the quantity μ_r which represents the maximum number of distinct out going arcs will be small. For example, when $\mathbf{a} \in CP$, $\mu_r = 1$ for all $r \in [0, n - 1]$.

While μ_r is more aptly a representative of the *degree of additive disorder*, we can view it as a representative of the *degree of additive structure* in an inverse way. Thus, the smaller is the value of μ_r the larger is the degree of additive structure and vice versa.

CLAIM 9.16. The configuration graph $\mathcal{C}$ includes both the combinatorial and relational aspects by capturing

- (i) the combinatorial aspect through the number of paths,
- (ii) the size m through segment lengths,
- (iii) and the degree of additive disorder/structure through μ_r .

9.4 Maximum number of paths in $\mathcal{C}$

PROPOSITION 9.17. Given a start point $(0, y_0) \in F_0$, the maximum number of distinct paths into level $r \in [n]$, denoted by ζ_r satisfies $\zeta_r < \mu_0^r$.

Proof. We feign ignorance regarding the edge in G_0 from which the local value y_0 comes from, and allow it from all edges of p_n . Hence the maximum number of distinct out going arcs from $(0, y_0)$ is at most μ_0 . Thereafter the maximum out degrees are at most $\mu_1, \mu_2, \dots$ and so on. Thus, the maximum number of distinct paths to level r is

$$\begin{aligned}\zeta_r &\leq \mu_0 \mu_1 \dots \mu_{r-1} \quad (\text{since some path lengths may coincide, we use } \leq) \\ &< \mu_0^r \quad (\text{since } \mu_r < \mu_0 \text{ for all } r \geq 1).\end{aligned}$$

$\square$

9.5 Some simple relations on γ_r and $|C(r, r + 1)|$

PROPOSITION 9.18. *Let r and $r + 1$ be two adjacent levels where $r \in [n - 1]$, with γ_r and γ_{r+1} points in each respectively. Then it is true that*

$$\gamma_{r+1} \leq \gamma_r \mu_r. \quad (9.1)$$

Proof. The inequality follows from the observation that the number of distinct points in level $r + 1$ cannot be larger than the number of distinct outgoing arcs emanating from all the distinct points of the previous level. $\square$

PROPOSITION 9.19. *Let r and $r + 1$ be two adjacent levels where $r \in [n - 1]$, with γ_r and γ_{r+1} points in each respectively. Then it is true that*

$$\max(\gamma_r, \gamma_{r+1}) \leq |C(r, r + 1)| \leq \gamma_r \gamma_{r+1}. \quad (9.2)$$

Proof. Since F_r as well as F_{r+1} are obtained from the paths in G_0 , it follows that these sets are connected by line segments. The minimum number of line segments needed to have γ_r points in level r , and γ_{r+1} points in level $r + 1$ is easily seen to be $\max(\gamma_r, \gamma_{r+1})$. If we have less than this number of segments, then at least one point will be isolated. The upper limit shows the other extreme, where every point in F_r is connected to every point in F_{r+1} . The actual value $|C(r, r + 1)|$ must necessarily be in between these limits. $\square$

PROPOSITION 9.20. *Let γ_r be the number of distinct points in level $r \in [0, n - 1]$. Then the maximum number of arcs between levels r and $r + 1$ is*

$$|C(r, r + 1)| \leq \gamma_r \mu_r \leq \gamma_r \mu_0. \quad (9.3)$$

Proof. Since the maximum out degree of a point $(r, y) \in F_r$ is at most μ_r , and there are γ_r distinct points in level r , the claimed inequality follows. $\square$

9.6 The importance of size

Recall that Ξ_n denotes the set of distinct configurations produced by the paths Π_n in the configuration graph $\mathcal{C}$.

PROPOSITION 9.21. *The maximum number of locally distinct configurations of paths in $\mathcal{C}$ is*

$$|\Xi_n| \leq \min(2^{mn}, \mu_0^n).$$

Proof. Since each segment s_r has 2^m available integral points, the maximum number of possibilities for distinct configurations is $2^m \times 2^m \times \dots \times 2^m$ (n times) which is equal to 2^{nm} . On the other hand, we know that the maximum number of distinct paths in $\mathcal{C}$ is bounded by μ_0^n . Since each distinct path results in a distinct configuration, putting them together, the claim follows. $\square$

Let n be fixed and consider instances of SSP with varying m . For each instance, we look at the corresponding configuration graph. If m is small, the segment lengths will be small and many paths will have the same configuration. If we use larger numbers i.e., increase m , the available number of configurations increase and some paths will start to have distinct configurations. As we continue to increase m , the number of distinct configurations in the geometry cannot increase arbitrarily since they cannot exceed the number of paths (see Section 13.4 for an example).

DEFINITION 9.22. For all $i \in [n]$, we define $\Gamma_i := \gamma_1 \gamma_2 \dots \gamma_i$.

COROLLARY 9.23. It is true that

$$|\Xi_n| \leq \min(\Gamma_n, \mu_0^n).$$

Proof. Whatever the sequence γ may be, it is clear that $\gamma_i \leq 2^m$ for all $i \in [n]$. Furthermore, it is easily seen that the maximum number of distinct configurations $|\Xi_n| \leq \gamma_1 \gamma_2 \dots \gamma_n$, and the claim follows. $\square$

While the number of distinct configurations cannot be increased by increasing m beyond $\log \mu_0$, the distribution of γ can be affected. Suppose that $m > \log \mu_0$. Then it can be seen that there are many distributions γ that support μ_0^n distinct configurations. For example,

$$\begin{aligned} \gamma &= (\mu_0, \mu_0, \dots, \mu_0) \text{ with } \max(\gamma) = \mu_0, \text{ and} \\ \gamma &= (\mu_0, \mu_0^2, \dots, \mu_0^n) \text{ with } \max(\gamma) = \mu_0^n, \end{aligned}$$

both support μ_0^n distinct configurations. In the first case each point in F_r connects with all points in F_{r+1} , whereas in the second case each point in F_r connects with μ_0 distinct points in F_{r+1} . The first case corresponds to *complete bi-partite layers* while the second corresponds to a μ_0 -regular tree.

Also while the number of distinct configurations cannot exceed μ_0^n , the above propositions don't tell us anything about $\max(\gamma)$. To obtain a bound for $\max(\gamma)$, our approach is to first obtain an upper bound for Γ_n , and then extract the desired upper bound. First of all, notice that Γ_n *need not be* bounded by μ_0^n . In fact it can well exceed it since the sequence γ can be wasteful, i.e., not all points in one level need to connect with all points in the next level. So the product Γ_n can be much larger than the number of paths.

9.7 A Product Inequality

In this section we derive a useful inequality, simply called a *Product Inequality* (PI), that would enable us to obtain an upper bound on $\max(\gamma)$.

DEFINITION 9.24. (Graph $\mathcal{J}$) Let the graph $\mathcal{J} := (V(\mathcal{J}), E(\mathcal{J}))$ be a layered graph with $n + 1$ levels, and n complete bipartite layers, defined by

$$\begin{aligned} V(\mathcal{J}) &:= V(\mathcal{C}) = \cup_{i=0}^n F_i, \\ E(\mathcal{J}) &:= \cup_{i=0}^{n-1} F_i \times F_{i+1}. \end{aligned}$$

Alternatively, $\mathcal{J}$ is a copy of the $\mathcal{C}$ graph, but where we add more arcs so that all points in level i are connected to all points in level $i + 1$ for all $i \in [0, n - 1]$. Thus, the number of arcs between levels i and $i + 1$ is simply $|F_i| \times |F_{i+1}| = \gamma_i \gamma_{i+1}$.

DEFINITION 9.25. (Arc Growth Factor (AGF) α_i) Let X be a layered graph with $l + 1$ levels and bipartite layers. Then for any $i \in [l]$ we define the arc growth factor $\alpha_i(X)$ as the ratio of the number of arcs between levels i and $i + 1$ to the number of arcs between levels $i - 1$ and i , i.e.,

$$\alpha_i(X) = \frac{\# \text{ arcs between levels } i \text{ and } i + 1}{\# \text{ arcs between levels } i - 1 \text{ and } i}.$$

PROPOSITION 9.26. For each $i \in [n - 1]$, the AGF in $\mathcal{J}$ satisfies

$$\alpha_i(\mathcal{J}) \leq \mu_{i-1} \mu_i < \mu_0^2. \quad (9.4)$$

Proof. By Definition 9.24 and by Definition 9.25, we have

$$\begin{aligned} \alpha_i(\mathcal{J}) &= \frac{\gamma_i \gamma_{i+1}}{\gamma_{i-1} \gamma_i} = \frac{\gamma_{i+1}}{\gamma_{i-1}} \\ &\leq (\mu_{i-1})(\mu_i) \quad (\text{using } \gamma_i \leq \gamma_{i-1} \mu_{i-1} \text{ and } \gamma_{i+1} \leq \gamma_i \mu_i) \\ &< \mu_0^2 \quad (\text{since } \mu_j < \mu_0 \text{ for all } j \in [n - 1]) \end{aligned}$$

and the claim follows. $\square$

REMARK 9.27. While the above Proposition 9.26 shows that $\alpha_i(\mathcal{J}) < \mu_0^2$, it is possible for $\alpha_i(\mathcal{J})$ to be less than one, for example, when $\gamma_{i+1} < \gamma_{i-1}$. This kind of a situation happens when paths fanout upto the middle levels, but later converge to a small number of points in the final level. In such cases $\alpha_i(\mathcal{J}) < 1$ for the latter levels.

The following proposition contrasts the graph $\mathcal{C}$ with $\mathcal{J}$ with regard to AGF.

PROPOSITION 9.28. For any $i \in [n - 1]$, the AGF in $\mathcal{C}$ satisfies

$$\alpha_i(\mathcal{C}) \leq \mu_i < \mu_0. \quad (9.5)$$

Proof. By the definition of AGF, we have

$$\begin{aligned}
\alpha_i(\mathcal{C}) &= \frac{|C(i, i+1)|}{|C(i-1, i)|} \\
&\leq \frac{\gamma_i \mu_i}{|C(i-1, i)|} \quad (\text{since } |C(i, i+1)| \leq \gamma_i \mu_i) \\
&\leq \frac{\gamma_i \mu_i}{\gamma_i} \quad (\text{since } |C(i-1, i)| \geq \max(\gamma_{i-1}, \gamma_i) \text{ by Equation (9.2)}) \\
&\leq \mu_i < \mu_0.
\end{aligned}$$

□

REMARK 9.29. From the above two propositions we see that $\alpha_i(\mathcal{C}) < \mu_0$ and $\alpha_i(\mathcal{J}) < \mu_0^2$. Thus, the maximum magnification factor per layer in $\mathcal{J}$ is a square of that of $\mathcal{C}$ to account for the presence of more arcs. In general, for any random instance the magnification factor in $\mathcal{J}$ is expected to be in between μ_0 and μ_0^2 .

PROPOSITION 9.30. Given a positive integer $r \in [n]$, let $\hat{\Upsilon}_r$ denote the number of distinct paths to level r in the $\mathcal{J}$ graph originating from the point $(0, y_0) \in F_0$ in the zeroth level. Then

$$\hat{\Upsilon}_r = \max_{j \in [r]} \gamma_1 \prod_{i=1}^{j-1} \alpha_i(\mathcal{J}).$$

Proof. For $r = 1$ the number of paths to level one is simply the number of arcs between levels zero and one, which is $\gamma_0 \gamma_1 = \gamma_1$ since $\gamma_0 = 1$. These paths result in γ_1 points from which new arcs fan out into the second level. The number of arcs between levels one and two, by definition is $\gamma_1 \gamma_2$. This can be likened to as the arcs to first level branching out to level two. Thus, the path magnification from level one to level two is $\frac{\gamma_1 \gamma_2}{\gamma_0 \gamma_1} = \alpha_1(\mathcal{J})$ by Definition 9.25. Thus, the number of paths to level two is $\gamma_1 \alpha_1(\mathcal{J})$ which equals $\gamma_1 \gamma_2$.

As we go to higher levels the arcs branch out into more arcs and we continue multiplying with AGFs $\alpha_i(\mathcal{J})$. However, by Remark 9.27, it is possible for some of the $\alpha_i(\mathcal{J})$ s to be less than one, particularly if all paths converge to a single point (or a small number of points) in level r . In this case, the number of distinct points in level r is small but the number of distinct paths is still the maximum obtained over all the previous levels. The $\max_{j \in [r]}$ term correctly preserves the maximum number if the maximum happens somewhere in the middle levels. □

COROLLARY 9.31. For $r \in [n]$, let $\hat{\Upsilon}_r$ denote the number of distinct paths to level r in the $\mathcal{J}$ graph originating from the point $(0, y_0) \in F_0$ in the zeroth level. Then $\hat{\Upsilon}_r < \mu_0^{2r-1}$.

Proof. We have

$$\begin{aligned}
\hat{\Upsilon}_r &= \max_{j \in [r]} \gamma_1 \prod_{i=1}^{j-1} \alpha_i(\mathcal{J}) \quad (\text{from Proposition 9.30}) \\
&< \gamma_1 (\mu_0^2)^{r-1} \quad (\text{from Proposition 9.26}) \\
&< \mu_0^{2r-1} \quad (\text{since } \gamma_1 \leq \mu_0)
\end{aligned}$$

and the claim follows. $\square$

Thus, the maximum number of paths in $\mathcal{J}$ is at most μ_0^{2r-1} which for algebraic simplicity and convenience (as will be seen later) is loosely bounded as μ_0^{2r} .

THEOREM 9.32. (*Product Inequality (PI)*) *Let $(\mathbf{a}, T)$ be an instance of the SSP with dimensions (n, m) , and let γ be the associated point number sequence. Then for any $r \in [n]$, it is true that $\Gamma_r < \mu_0^{2r}$.*

Proof. The number of distinct paths in $\mathcal{J}$ is simply the number of distinct combinations taking one point from each of the sets F_i for $i \in [r]$, which is clearly $\Gamma_r = \gamma_0 \gamma_1 \dots \gamma_r$. On the other hand we know from Corollary 9.31 that the number of distinct paths is less than μ_0^{2r} . Putting them together the claim follows. $\square$

9.8 An upper bound on $\max(\gamma)$

In this section we obtain an upper bound on $\max(\gamma)$ for any given instance of the SSP. We also obtain an asymptotic formula for the profile of γ in the absence of any additive structure and for large m . These results depend crucially on the *Product Inequality*.

Taking the r -th root on both sides of $\Gamma_r < \mu_0^{2r}$, we obtain that the geometric mean $\gamma_{GM} := \Gamma_r^{1/r} < \mu_0^2$. Thus, the PI simply says that the geometric mean γ_{GM} is bounded by a polynomial. However, it doesn't immediately reveal anything about $\max(\gamma)$ since one can envision an arbitrary number of sequences γ with varying maximal values that all have a given bounded geometric mean. On the other hand, the sequence γ cannot be arbitrary as it has to correspond to paths on the OL and satisfy the growth condition given in Equation (9.1). So what is the worst case distribution that results in maximum possible $\max(\gamma)$? A momentary thought shows that the worst case must correspond to a Geometrical Progression.

THEOREM 9.33. *For any instance of SSP of size n , we have $\max(\gamma) < \mu_0^4$.*

Proof. To maximize $\max(\gamma)$ we consider the sequence $(\gamma_1, \gamma_2, \dots, \gamma_r)$ to be a $GP(r, \theta)$ for some $1 < \theta \leq \mu_0$. In particular, we let $\gamma_i = \theta^i$ for $i \in [k]$. Using the product

inequality (Theorem 9.32) we have $\Gamma_r < \mu_0^{2r}$. Thus,

$$\begin{aligned}\theta^1 \theta^2 \dots \theta^r &< \mu_0^{2r}, \\ \theta^{\frac{r(r+1)}{2}} &< \mu_0^{2r}, \\ \theta^{\frac{(r+1)}{2}} &< \mu_0^2, \text{ (after taking the } r \text{ th root on both sides)} \\ \theta^{(r+1)} &< \mu_0^4, \text{ (after squaring on both sides)} \\ \theta^r &< \frac{\mu_0^4}{\theta} < \mu_0^4, \text{ since } \theta > 1.\end{aligned}$$

Thus, the maximum element $\max(\gamma) = \gamma_r < \mu_0^4$. $\square$

It should be noted that this is an extreme case that cannot be realized in practice since μ_r is decreasing function of r , but does provide us with a (loose) upper bound for $\max(\gamma)$. Next we show that the above result is still true even if we consider a symmetric geometric progression.

DEFINITION 9.34. (*Symmetric Geometric Progression SymGP(r, θ)*) Let r be an even positive integer, and let

$$\text{SymGP}(r, \theta) := (\theta^1, \theta^2, \dots, \theta^{r/2}, \theta^{r/2}, \theta^{r/2-1}, \dots, \theta^1)$$

be a r -term Symmetric Geometric Progression centered at $r/2$, where $\theta > 1$.

PROPOSITION 9.35. If γ is a SymGP(r, θ) where $\theta > 1$, and $\gamma_{GM} < \mu_0^2$, then $\max(\gamma) \leq \mu_0^4$.

Proof. For simplicity we assume that r is even. We have

$$\begin{aligned}\Gamma_r &= (\theta^{\frac{(r/2)(r/2+1)}{2}})^2, \\ &= \theta^{(r/2)(r/2+1)} = \theta^{\frac{r(r+2)}{4}}\end{aligned}$$

$$\gamma_{GM} = \Gamma_r^{1/r} = \theta^{\frac{r+2}{4}} < \mu_0^2.$$

Squaring both sides, we get

$$\begin{aligned}\theta^{r/2+1} &< \mu_0^4, \\ \text{which implies } \theta^{r/2} &< \frac{\mu_0^4}{\theta} < \mu_0^4\end{aligned}$$

since $\theta > 1$. Thus, $\max(\gamma) < \mu_0^4$. $\square$

Note that even if r in the above proposition is odd the result will be unaffected.

REMARK 9.36. We have seen above that for a Geometric Progression as well as a Symmetric Geometric Progression, we have $\max(\gamma) < \mu_0^4$. Even if we use any of the following forms,

- (i) A Geometric Progression followed by a flat region,
- (ii) A Geometric Progression followed by a flat region that is followed by a decreasing Geometric Progression,

we will still get the same bound of μ_0^4 . Since the proofs simply mimic the proof of Theorem 9.33, they are not given here.

PROPOSITION 9.37. *For any given instance of SSP, if we use Algorithm 4, then the peak growth factor satisfies $\eta_{peak} < \mu_0^4$.*

Proof. (By contradiction) Suppose that $\eta_k > \mu_0^4$ at some iteration $k \in [m]$ where $k > 4 \log \mu_0$. Then there must be at least one edge $e_i \in G_0$ with more than μ_0^4 disjoint parts that survived the FILTER steps. Since the number of distinct valid paths is less than μ_0^4 by Theorem 9.33, we have more parts than distinct paths. By the Pigeon Hole Principle, there must be at least one part of e_i , with no valid paths through it. Such an edge will either be above or below the OL but does not intersect it. In other words, either the shortest path is larger than the edge length, or the longest path is negative. In either case the FILTER step (see Algorithm 3, step 9) removes the node by Proposition 8.12 providing a contradiction, and the claim follows. $\square$

The above bound will be significantly improved later when we analyze the Zero Paths. For now, we already have a polynomial bound of $\mu_0^4 = \frac{n^{12}}{81}$ (from Definition 9.14), and thus a polynomial time complexity for Algorithm 4.

9.9 Profile of γ

Given an instance of the SSP, what is the shape or profile of $\gamma : [n] \rightarrow \mathbb{N}$ as a function of level? The short answer is that it depends on n , m and the *additive structure* of the given instance. For instance, if $m \leq \log \mu_0$, then all the segments in the graph $\mathcal{C}$ have a length at most μ_0 , and the profile of γ in the worst case is capped by a uniform distribution with a maximum of μ_0 . For larger m , if there is a high degree of additive structure, as in when $\mathbf{a} \in CP$ or $\mathbf{a} \in AP$, then many arc weights will be the same and μ_r will be very small. Thus, the number of distinct configurations in $\mathcal{C}$ will be small, and hence $\max(\gamma)$ will also be small.

For the case when $m > \log \mu_0$, and when there is no additive structure, the profile of γ is a Gaussian with peak at $r = n/2$. This should not be surprising, since after all the number of locally distinct paths to level r is a subset of the number of transformation paths to level r , which as we have already seen follows a Binomial distribution (see Theorem 5.6). However, we also noted earlier that only a subset of the transformation paths result in valid wormhole paths for a given target value T , and even a smaller number of distinct OLs. In other words, we have the sequence of inequalities

$$\gamma_r \leq \hat{\beta}_T(r) \leq \beta_T(r) \leq \beta(r) = \binom{n}{r}.$$

While γ_r is loosely bounded by the combinatorial quantity $\binom{n}{r}$ on one side, it is not determined by it. The *main controlling factor* is the number of unique arc weights μ_r that reflects the degree of additive structure. As we have noted before, the simplest proof of this is furnished by the $\mathbf{a} \in CP$ case, where $\mu_r = 1$ and $\max(\gamma) = 1$. We will see other cases in Section 11.

Thus, for any given instance of SSP with sufficiently large m and no additive structure we have

$$\gamma_r = K e^{-\frac{(r-n/2)^2}{2\sigma^2}}, \forall r \in [n],$$

where $K = K(n)$ is a positive scaling factor, and σ^2 is the variance. We will determine bounds on K and σ^2 below which depend on μ_0 .

For intermediate values of m and differing degrees of additive structure, the profile of γ can take several different shapes, still bounded by Gaussian distribution from above.

9.10 Asymptotic profile of γ

Consider an instance $(\mathbf{a}, T)$ with large m and large μ_r . When m is large, say $m > n$, the segment lengths in $\mathcal{C}$ are large which allows for the possibility to have more distinct points in each segment. When μ_r is large we have larger gain in the configuration graph allowing for more distinct paths. This is the case of primary interest, where we are interested in the profile of γ as $n \rightarrow \infty$.

DEFINITION 9.38. (*Gaussian Distribution* $Gauss(n, n/2, \sigma^2)$) For a positive integer n , let

$$f(i) = K e^{-\frac{(i-n/2)^2}{2\sigma^2}}, \text{ for } i \in [n], \quad (9.6)$$

be a Gaussian distribution with scaling factor $K > 0$, mean $n/2$ and variance σ^2 .

PROPOSITION 9.39. Let $\gamma \sim Gauss(n, n/2, \sigma^2)$ be a Gaussian distribution with scaling factor K . Then $K < \mu_0^3$ and $\sigma^2 = \frac{n^2}{8 \log K}$.

Proof. Taking log on both sides of Equation (9.6), we have

$$\ln(\gamma_i) = \ln K - \frac{(i - n/2)^2}{2\sigma^2}.$$

Summing over $i \in [n]$ on both sides,

$$\begin{aligned} \sum_{i=1}^n \ln(\gamma_i) &= n \ln K - \sum_{i=1}^n \frac{(i - n/2)^2}{2\sigma^2}, \\ &= n \ln K - \frac{1}{2\sigma^2} \frac{2n + n^3}{12}, \\ &= n \ln K - \left(\frac{n(n^2 + 2)}{24\sigma^2} \right). \end{aligned}$$

Applying the product inequality, we can write

$$\begin{aligned} n \ln K - \left(\frac{n(n^2 + 2)}{24\sigma^2} \right) &< 2n \ln \mu_0, \\ \ln K - \left(\frac{(n^2 + 2)}{24\sigma^2} \right) &< 2 \ln \mu_0, \\ \ln K &< 2 \ln \mu_0 + \left(\frac{(n^2 + 2)}{24\sigma^2} \right), \end{aligned}$$

leading to

$$K < \mu_0^2 e^{\frac{(n^2+2)}{24\sigma^2}}. \quad (9.7)$$

To determine the variance we use the initial value of $\gamma_0 = 1$, i.e.,

$$\begin{aligned} 1 = \gamma_0 &= K e^{-\frac{(-n/2)^2}{2\sigma^2}}, \\ \implies \frac{1}{K} &= e^{-\frac{n^2}{8\sigma^2}}. \end{aligned}$$

Taking log on both sides and simplifying, we get

$$2\sigma^2 = \frac{n^2}{4 \log K}.$$

Using this in Equation (9.7), we get

$$K < \mu_0^2 e^{\frac{(n^2+2)}{24\sigma^2}} = \mu_0^2 e^{\frac{n^2+2}{12(\frac{n^2}{4 \log K})}}.$$

After simplification, we get

$$K < \mu_0^2 e^{(\frac{n^2+2}{n^2}) \frac{\log K}{3}} \approx \mu_0^2 e^{\frac{\log K}{3}}, \quad (\text{for large } n).$$

Taking log on both sides, we get

$$\begin{aligned} \log K &< 2 \log \mu_0 + \frac{\log K}{3}, \\ \text{and } K &< \mu_0^3. \end{aligned}$$

□

Notice that this bound of $\max(\gamma) < \mu_0^3$ is better than what we obtained for a geometric progression above.

EXAMPLE 9.40. The asymptotic distribution of γ for a random $n = 36, m = 72$ instance and a random $n = 40, m = 80$ instance with no additive structure is shown below in Figure 16.




Figure 16: The top plot is for a $n = 36, m = 72$ example, where the profile of γ is shown by red dots. The blue curve corresponds to a Gaussian fit where we used μ_0^δ instead of μ_0^2 as the gain in $\mathcal{J}$ graph (see Remark 9.27 and Remark 9.29), where $\delta = 1.2225$ and $\mu_0 = \frac{n^3+3n^2-13n+6}{3} = 16,694$. The bottom plot is for a $n = 40, m = 80$ example, where we used $\delta = 1.385$ with $\mu_0 = \frac{n^3+3n^2-13n+6}{3} = 22,762$. As n increases, we expect $\delta \rightarrow 2$.

THEOREM 9.41. *For an instance of SSP without any additive structure and with large m , the associated sequence γ satisfies the asymptotic profile given by*

$$\gamma_r = \mu_0^3 e^{-\frac{6(\log \mu_0^2)(r-n/2)^2}{n^2}}, \quad \forall r \in [n], \quad (9.8)$$

and $\gamma_0 = 1$.

COROLLARY 9.42. *For any instance of SSP with dimensions (n, m) , no edge equivalence class can intersect the orbital line at more than μ_0^3 distinct points.*

The above corollary suggests that there is an *occupancy limit* in the segments of $\mathcal{C}$ graph!

9.11 Geometric Interpretation

We provide a rough geometrical interpretation on what the above result means, in this section. However, it is only a qualitative picture aimed at complementing the rigorous quantitative answer given above.

DEFINITION 9.43. ($\text{ken}((T - M, T])$) *Consider the geometric structure as in Figure 8 and assume that we have a window along the y -axis with limits $(T - M, T]$. We further assume that all the lower vertices of edges that intersect the OL are lighted up, and the rest of the points in the geometry are dark. When we look through this window horizontally, the number of distinct lighted points we see is called $\text{ken}((T - M, T])$, which is nothing but the projection of these points on to the y axis.*

If the given instance $\mathbf{a} \in CP$, then it is clear that $\text{ken}((T - M, T]) = 1$ since all edges that intersect the OL do so at their lower vertex if there is a solution, otherwise all of them will be offset by the same quantity (see Figure 18). Thus, there is only one lighted point we can see through the window. If $\mathbf{a} \in AP$, then $\text{ken}((T - M, T]) \leq n$ which is proved in Section 11. For a random instance the above result suggests that $\text{ken}((T - M, T]) \leq n\mu_0^3$ since there are n levels in the configuration graph. To understand this intuitively, recall from Section 5.4 that there are many segments which have an exponential path multiplicity, i.e, an exponential number of NDPs pass through them. Also note that in the case when all elements of $\mathbf{a}$ are distinct, the graph of NDPs slowly rises from left to right, and many subspaces go out of range with respect to a fixed OL. This becomes clear when we consider the more drastic example of $\mathbf{a} = \mathbf{b}$. In this case only one of either the left or right half subspaces intersect the given OL, but not both. To summarize, when there is lot of additive structure many points have the same projection; When there is no additive structure the number of distinct intersections is fewer because of the rise in the graph, and segments having an exponential path multiplicity.

9.12 The case of Zero Paths

So far we have been looking at the case of all valid paths on the OL. However, in the SSP we are primarily interested in paths on the OL that lead to a solution, which is

clearly a subset of all valid paths on the OL. While we could simply use the bounds obtained for *all valid paths*, there is benefit in considering the case of *all zero paths* since it results in a significant improvement in the bound for $\max(\gamma)$.

In this section, we investigate the configuration graph where every path is a *zero path*. Recall that a *zero path* is a valid path that ends at the vertex point of an edge in G_0 , or equivalently the point $(j, 0)$ in the $\mathcal{C}$ for some $j \in [n]$. When only zero paths are plotted in the configuration graph, we call it as the ZP-configuration graph.




Figure 17: A ZP-Configuration graph consisting of zero paths. Only some sample paths are shown. The incoming paths into lower end of a segment from the previous segment are shown in green. The other paths are shown in red.

In Figure 17 we show a rough sketch of the structure of ZeroPaths on the OL. Every path must end at $(r, 0)$ for some $r \in [n]$.

In what follows and for the rest of this subsection we assume that we have *only zero paths* in the configuration graph. As in the case of all valid paths we use F_r , with $r \in [n]$ to denote the set of distinct points in each segment.

PROPOSITION 9.44. *For a positive integer $r \in [n]$, let $(r, y) \in F_r$ be a point in segment s_r . Then either $y = 0$ or there is at least one zero path from (r, y) to some $(k, 0)$ where $k \in \{r + 1, r + 2, \dots, n\}$.*

Since any zero path must end at a vertex point, the above proposition is clear. Instead of considering zero paths of all lengths, we first consider zero paths of a fixed length k for some $k \in [n]$ for convenience. Later we will generalize to the case of all lengths.

DEFINITION 9.45. (*Zero Paths of length k*) *For a given positive integer $k \in [n]$, let ZP_k denote the set of all zero paths of length equal to k . The associated point number sequence is denoted by*

$$(\dot{\gamma}(k, 0), \dot{\gamma}(k, 1), \dots, \dot{\gamma}(k, k))$$

where $\dot{\gamma}(k, 0) = \gamma_0 = 1$. For the sequence we use double index, where the first index corresponds to path length and second the level number. Since all these paths end at $(k, 0) \in F_k$ and we have only zero paths, it follows that $\dot{\gamma}(k, k) = 1$. The set of zero paths ZP_k have a common source point $(0, y_0) \in s_0$ and a common destination point $(k, 0) \in s_k$. Let $|ZP_k|$ denote the number of zero paths of length k .

PROPOSITION 9.46. *The sequence $(\dot{\gamma}(k, 1), \dot{\gamma}(k, 2), \dots, \dot{\gamma}(k, k))$ satisfies*

(i) $\dot{\gamma}(k, r+1) \leq \dot{\gamma}(k, r)\mu_r \leq \dot{\gamma}(k, r)\mu_0$, for all $r \in [k-1]$

(ii) $\dot{\gamma}(k, 0) = 1$ and $\dot{\gamma}(k, k) = 1$.

Proof. Since the set ZP_k is a subset of all valid paths, it follows that the sequence $(\dot{\gamma}_0, \dots, \dot{\gamma}_k)$ obeys the same growth bounds as in Equation (9.1). This takes care of (i). Item (ii) directly follows from Definition 9.45. $\square$

PROPOSITION 9.47. *The number of distinct zero paths of length k is at most $\mu_0^{k/2}$.*

Proof. Consider the segment s_r where $r \in [k-1]$ is an intermediate level in the ZP-configuration graph. Then the number of paths into s_r from the zeroth level is at most $\mu_0\mu_1 \dots \mu_{r-1} \leq \mu_0^r$. The number of (reverse) paths into s_r from the point $(k, 0) \in s_k$ is at most $\mu_k\mu_{k-1} \dots \mu_{r+1} \leq (\mu_0)^{k-r}$.

Since we are considering only zero paths, it follows that the each path from the left must match with a path from the right; otherwise we would have non-zero paths. Thus,

$$\# \text{ distinct zero paths through } s_r \leq \min(\mu_0^r, \mu_0^{k-r}).$$

To maximize the number of distinct zero paths through s_r we equate both to obtain that $r = k/2$. Thus, the maximum number of distinct zero paths of length k is at most $\mu_0^{k/2}$. $\square$

LEMMA 9.48. (*Product Inequality for Zero Paths*) *Consider the configuration graph of the set of all zero paths of length $k \in [n]$ with source point $(0, y_0) \in F_0$ and destination point $(k, 0) \in F_k$. Let the associated point number sequence be $\dot{\gamma}_k := (\dot{\gamma}(k, 0), \dots, \dot{\gamma}(k, k))$ and let $\dot{\Gamma}_k := \dot{\gamma}(k, 0) \dots \dot{\gamma}(k, k)$. Then $\dot{\Gamma}_k < \mu_0^k$.*

Proof. We will mimic the proof as in Corollary 9.31, by first treating the product $\dot{\Gamma}_k = \dot{\gamma}(k, 0) \dots \dot{\gamma}(k, k)$ as the number of distinct paths in a configuration graph (denoted by $\mathcal{J}$) where each point in a level is connected to all points in the neighboring levels. Furthermore, we have the arc growth factors (AGF defined in Definition 9.25) $\alpha_1, \alpha_2, \dots, \alpha_{k-1}$. Since the sequence $\dot{\gamma}$ follows the same growth bounds as the parent sequence γ (the sequence corresponding to all valid paths), it follows from Proposition 9.26 that $\alpha_i < \mu_0^2$ for all $i \in [k/2 - 1]$. For levels greater than or equal to $k/2$ the AGFs will be less than one since the paths converge to a single point in level k . We therefore set

$$\alpha_i = \begin{cases} \mu_0^2 & \text{for all } i \in [k/2 - 1], \\ 1 & \text{for all } i \in \{k/2, k/2 + 1, \dots, k-1\}. \end{cases}$$

The number of arcs between the zeroth and first levels is at most μ_0 . With this the number of paths in $\mathcal{J}$ can be estimated using AGFs as

$$\# \text{ paths in } \mathcal{J} \leq \mu_0(\mu_0^2)^{k/2-1} = \mu_0^{k-1} < \mu_0^k.$$

Since $\dot{\Gamma}_k$ represents all possible combinations (paths) it follows that

$$\dot{\Gamma}_k < \mu_0^k.$$

$\square$

LEMMA 9.49. Consider the set of all zero paths of length k , with associated point number sequence $\hat{\gamma}_k := (\hat{\gamma}(k, 0), \dots, \hat{\gamma}(k, k))$. Then $\max(\hat{\gamma}_k) < \mu_0^2$.

Proof. To maximize $\hat{\gamma}(k, k/2)$ we consider the sequence $(\hat{\gamma}(k, 1), \dots, \hat{\gamma}(k, k))$ to be a $SymGP(k, \theta)$ for some $1 < \theta \leq \mu_0$. In particular we let $\hat{\gamma}(k, i) = \theta^{\min(i, k-i)}$ for $i \in [k]$. From Lemma 9.48 we have $\hat{\Gamma}_k < \mu_0^k$. Thus,

$$\begin{aligned} (\theta^1 \theta^2 \dots \theta^{k/2})^2 &< \mu_0^k, \\ \theta^{\frac{k(k+2)}{4}} &< \mu_0^k, \\ \theta^{\frac{k+2}{2}} &< \mu_0^2, \text{ (after taking the } k/2 \text{ th root on both sides)} \\ \theta^{k/2+1} &< \mu_0^2, \\ \theta^{k/2} &< \frac{\mu_0^2}{\theta} < \mu_0^2, \text{ since } \theta > 1. \end{aligned}$$

Thus, the maximum element $\hat{\gamma}(k, k/2) < \mu_0^2$. $\square$

THEOREM 9.50. Consider the set of all zero paths of all lengths $k \in [n]$ for a given instance of the SSP, denoted by ZP_* . Let the associated point number sequence be $\bar{\gamma} := (\bar{\gamma}_0, \bar{\gamma}_1, \bar{\gamma}_2, \dots, \bar{\gamma}_n)$, where $\bar{\gamma}_0 = 1$. Then

$$\max(\bar{\gamma}) < n\mu_0^2.$$

Proof. It is clear that

$$ZP_* = \cup_{k=1}^n ZP_k.$$

Since for each fixed $k \in [n]$, the maximum number of distinct points in $s_{k/2}$ is less than μ_0^2 , it follows that the number of distinct points in any segment will be at most n times μ_0^2 and the claim follows. $\square$

The following theorem follows from the above result.

THEOREM 9.51. The maximum number of distinct zero paths through any node $e_i \in G_0$, is at most $\min(2^m, n\mu_0^2)$.

Note that we cannot claim a Gaussian profile for γ in the case of zero paths since there may not even be a solution to a given instance. As a result, we used the symmetric geometric progression to get a worst case estimate in Lemma 9.49. Since we are bounding the number of distinct zero paths into any node using the above result (which includes contributions from all nodes in a given level represented in a single segment of the configuration graph), it should be clear that the above bound is not tight.

10 Main Results & Complexity Calculations

In this section we prove the main claims of the paper, which follow easily given all the results up to this point. Recall that the complexity of the algorithm depends on the bound for η_{peak} (Proposition 8.17). We have already proved a bound in Proposition 9.37 using the bound for maximum number of valid paths through a node. We provide a much better bound next using the results of our analysis on distinct zero paths.

THEOREM 10.1. $\eta_{peak} \leq \min(2^m, \frac{n^7}{9})$.

Proof. (Sketch) We have already shown that $\eta_{peak} < \mu_0^4$ (in Proposition 9.37) using the bound for the maximum number of distinct valid paths through any node. This bound is too loose for η_{peak} , since the FILTER step of Algorithm 4 not only removes nodes that have only invalid paths into them, but also removes nodes that have valid but only non-zero paths. Only those nodes that can support zero paths are preserved by the FILTER step. To show the claim we make the following arguments.

There are three types of paths in the orbital graph: (1) Invalid paths, (2) Valid but non-zero paths and (3) Zero paths. For an arbitrary node $e_i \in G_0$, all the three types of paths can co-exist and go through the node. We know from Theorem 9.51 that the maximum number of distinct zero paths cannot exceed $n\mu_0^2$. Thus, after $k := \lceil \log(n\mu_0^2) \rceil$ refinements, we can expect, in general, that all the distinct zero paths are separated into the 2^k different parts of e_i . Each part may have paths of types (1) and (2) along with the zero path. Let e_j be one such part of e_i , that has all the three types of paths. In the next refinement, the node e_j is split into e'_j and e''_j . The single distinct zero path can only be in one of e'_j or e''_j but not both. The other part will now have only paths of types (1) or (2). If it is of type (1), the node is removed since it does not intersect the OL. If it is of type (2), the node is either already labeled FALSE or will now be labeled FALSE. The FILTER routine disconnects this node from the destination node e_∞ , and will be removed when all the FALSE LEAF nodes are removed during the n loops of the FILTER procedure (see Remark 8.15). It is possible for a few FALSE nodes that do not support zero paths to survive the current FILTER step due to the way they are connected to next level, but the uncertainty will be removed in one or two refinements. These small fluctuations do not affect, since one can see that the nodes in the first few levels of the orbital graph will have only a very small number of surviving parts (and they cannot increase with more refinements), as the number of paths into them is very small. In other words, whatever the amount by which the number of nodes increases due to the REFINE step beyond iteration k , about the same number get removed in general by the FILTER procedure, since the distinct zero paths have been separated. Thus, the growth factor of the graph cannot increase beyond $n\mu_0^2 = \frac{n^7}{9}$ (from Definition 9.14). Taking into account the edge length limit of 2^m also, the claim follows. $\square$

COROLLARY 10.2. *For any positive number n , there is a positive number $m_0 = \lceil 7 \log n \rceil$ such that for all $m > m_0$, size of elements in $\mathbf{a}$ do not have any affect on the complexity of the SSP.*

COROLLARY 10.3. *The maximum number of nodes in the refined graph G_m is at most $\frac{n^{11}}{216}$. The total number of arcs in G_m is at most $\frac{n^{13}}{6480}$.*

Proof. The maximum number of nodes in the refined graph G_m is,

$$|V(G_m)| \leq \eta_{peak}|V_0| = \min(2^m, \frac{n^7}{9}) \frac{n^4}{24} \leq \frac{n^{11}}{216}.$$

Every arc in the initial graph G_0 can become at most η_{peak} arcs in G_m . Thus,

$$|E(G_m)| \leq \eta_{peak}|E_0| = \min(2^m, \frac{n^7}{9}) \frac{n^6}{720} \leq \frac{n^{13}}{6480}.$$

□

With this we finally have the main result:

THEOREM 10.4. $\mathcal{P} = \mathcal{NP}$.

From a result of Stockmeyer [26], the following corollary follows, where $\mathcal{PH}$ is the polynomial hierarchy (the set of classes above $\mathcal{NP}$ and $co\text{-}\mathcal{NP}$ with increasing number of alternating quantifiers $\exists$ and $\forall$).

COROLLARY 10.5. $\mathcal{P} = \mathcal{PH}$.

10.1 Finding the number of solutions

With our algorithm, finding the number of solutions to a given instance is easy, and practically free. Once we obtain the *Solution Graph* G_m , we use a possibly well-known trick from Graph Theory, to find the number of solutions N_{sols} . If $G_m = \emptyset$, then $N_{sols} = 0$. Otherwise, we follow these steps below.

- (i) Attach a counter ξ_i to each node $e_i \in G_m$.
- (ii) Set $\xi_0 = 1$ for the root node.
- (iii) For each node $e_i \in G_m$, let

$$\xi_i = \sum_{j \in \text{parents}(e_i)} \xi_j.$$

- (iv) Traverse the graph using a Breadth-First-Search (BFS), applying (iii) at each node.
- (v) Recall that the node e_∞ is connected to all the *TRUE* nodes in the graph, and that it acts as a collector. Thus, the number of solutions N_{sols} to the given instance is simply

$$N_{sols} = \xi_\infty.$$

Since G_m is polynomial sized, the complexity of running the above steps is simply $O(|V_m| + |E_m|)$, which is a polynomial. This proves the third claim of this paper.

THEOREM 10.6. $\mathcal{FP} = \#\mathcal{P}$.

Since $\mathcal{P}$ with a $\#\mathcal{P}$ oracle is simply $\mathcal{P}$ by the above theorem, we have that

COROLLARY 10.7. $\mathcal{P} = \mathcal{P}^{\#\mathcal{P}}$.

Bernstein and Vazirani [4], defined the complexity class $\mathcal{BQP}$ of problems that can be solved in polynomial time using quantum computers. They further showed the containment that $\mathcal{BQP} \subseteq \mathcal{P}^{\#\mathcal{P}}$. Therefore the following (somewhat surprising) result readily follows.

COROLLARY 10.8. $\mathcal{BQP} \subseteq \mathcal{P}$.

10.2 Classifying the solutions

The layered structure of the graph G_0 (or G_m), with a unique 5-tuple identification for the nodes, provides a natural classification of the solutions. This classification may be used possibly in several ways, for example: It is quite conceivable that in some problem (that could be reduced to *Subset Sum Problem*), additional requirements among equivalent solutions, could be important. Suppose that there are two solutions, with the first zero path terminating at $e_i \in G_0$, and the second terminating at $e_j \in G_0$. While both solutions solve the original problem, one solution may be *less expensive or easier* to implement than the other. Mapping the solution classes back to the original problem will provide the needed information. A second application is where one is interested in a particular equivalence class (e.g in level 10, on a q_7 curve, etc.), and whether there are any solutions into it. In such a case, we can consider the smaller graph reachable from e_0 in 10 hops with a fixed destination node, and compute paths.

10.3 Numerical results

We provided several numerical examples to illustrate the workings of the algorithm, and also point out some applications of the algorithm to some combinatorial problems. These are put at the end of the paper (see Section 13) so as to maintain the flow of the paper. We urge the reader to take a look at it, as it helps to understand the algorithm better.

11 Instances & Additive Structure

We have seen that given any instance, we can find the number of solutions using a polynomial amount of resources. This raises the following questions:

QUESTION 11.1. *Does every instance $(\mathbf{a}, T)$ have the same complexity? If not, then what controls the complexity?*

We have already shown that the *degree of additive disorder* (equivalently *degree of additive structure* in an inverse way) of $\mathbf{a}$, as measured by the number of unique arc weights μ_r where $r \in [n]$, is the main determiner of the *complexity*. We have shown that when μ_r is small, such as in the case of $\mathbf{a} \in CP$, the complexity is low. In this section we consider more commonly used measures of *additive structure* given below. Then we consider some instances with varying degrees of additive structure and provide the complexity estimates, while also attempting to make a connection between the *additive structure* and *geometric structure*.

11.1 Two quantities of interest

There does not seem to be a standard definition for the *degree of additive structure*, but several possibilities are listed in the book by Tao and Vu [27]. To understand the degree of *additive structure*, we consider two quantities mentioned in [27], that are appropriate for the *Subset Sum Problem*.

DEFINITION 11.2. (*Unique sums*) Given all the 2^n subset sums for a given $\mathbf{a}$, we denote the set of all unique sums by $\mathcal{U}$, and the number of unique sums by $U = |\mathcal{U}|$. Clearly $1 \leq U \leq \min(2^n, A_n)$.

DEFINITION 11.3. (*Maximum Concentration*) Given all the 2^n subset sums for a given $\mathbf{a}$, the maximum concentration denoted by N_{LO} is defined as

$$N_{LO} = \max_{s \in \mathcal{U}} |\{\mathbf{x} \in \{0, 1\}^n : \mathbf{a} \cdot \mathbf{x} = s\}|$$

i.e., the maximum number of subsets that have the same sum. This type of quantity was first studied by Littlewood and Offord in what is now called Littlewood-Offord Theory [27], and hence the subscript LO .

It would be useful to generalize the notion of *Maximum Concentration* to simply *concentration* for each unique sum $s \in \mathcal{U}$. For this, we order the unique sums in increasing order, and denote and define L_i as the set of all subsets that have the same i th unique sum. Formally,

$$L_i := \{\mathbf{x} \in \{0, 1\}^n : \mathbf{a} \cdot \mathbf{x} = s_i, s_i \in \mathcal{U}\}$$

where s_i is the i th unique sum. Thus, $N_{LO} = \max_i |L_i|$. Clearly, we have

$$\sum_{i=1}^U |L_i| = 2^n.$$

Note that $\mathcal{L} := \{|L_1|, |L_2|, \dots, |L_{U-1}|, |L_U|\}$ is nothing but the distribution of sums in $[0, A_n]$ into U bins, and this captures more information about the point set compared to simply N_{LO} .

The quantities U and N_{LO} are inversely proportional. If U is large, N_{LO} will be small and vice versa. When there is additive structure, $|L_i|$ s are expected to be large and U is expected to be small. Given an instance $\mathbf{a}$, the subset sums will be distributed

according to L . For a given T , the only possible number of solutions is from the set $\{0, |L_1|, |L_2|, \dots, |L_{U-1}|, |L_U|\}$ and nothing else. For a given instance $(\mathbf{a}, T)$, suppose that there are $|L_i|$ solutions for some $i \in [U]$. What is the relation between all these subsets apart from being solutions? One ready answer is provided by the *Solution Graph*, where all of these are connected to a common root node.

11.2 Dependence of Solution Graph size on U

Next, we show how the size of the final graph, in particular $|V_m|$ can be related to the number of unique sums U . The geometric structure we discussed earlier enables this relation. Although it is a rather loose bound, it sheds light on how the *additive structure* (via U) affects the complexity.

LEMMA 11.4. $|V_m| \leq \min(\frac{2U}{n(n+1)}, 2^m) |V_0|$.

Proof. Let ζ_k be the number of unit intervals of a node $e_k \in V_0$ that appear in the final graph. Then it can be seen that

$$|V_m| = \sum_{k=1}^{|V_0|} \zeta_k.$$

Note that for each node e_k , we have $|e_k| \leq 2^m$. Since each edge is on a non-decreasing path from $(0, 0)$ to (B_n, A_n) with $n(n+1)/2$ points, every unique intersection point with the OL contributes $n(n+1)/2$ points (i.e., the points on the curve) to the count U of unique sums. Since there are U unique sums, it follows that $\zeta_k n(n+1)/2 \leq U$. Alternatively,

$$\zeta_k \leq \frac{2U}{n(n+1)}.$$

Using the fact that $\zeta_k \leq |e_k|$, we get

$$\zeta_k \leq \min(\frac{2U}{n(n+1)}, 2^m).$$

Thus,

$$|V_m| = \sum_{k=1}^{|V_0|} \zeta_k \leq \sum_{k=1}^{|V_0|} \min(\frac{2U}{n(n+1)}, 2^m) = \min(\frac{2U}{n(n+1)}, 2^m) |V_0|,$$

and the claim is proved. $\square$

COROLLARY 11.5. *For a given instance $(\mathbf{a}, T)$, if the number of unique sums U is bounded by a polynomial, then the size of the solution graph, i.e., $|G_m|$ is also bounded by a polynomial.*

COROLLARY 11.6. *Let $\mathbf{a} \in CP$. Then given $(\mathbf{a}, T)$, we have $|V_m| \leq \frac{2|V_0|}{n}$.*

Since $U = n + 1$ when $\mathbf{a} \in CP$, the above corollary follows from Lemma 11.4.

PROPOSITION 11.7. Let $\mathbf{a} = AP(n, c, d)$. Then the number of unique sums U in the solution space is at most $\frac{n^3+5n+6}{6}$.

Proof. We have $a = \{c + (n-1)d, c + (n-2)d, \dots, c + 2d, c + d, c\}$. A sum with k elements has a sum in the range $[kc + \frac{(k-1)k}{2}d \quad kc + (\frac{(n-1)n}{2} - \frac{(n-k-1)(n-k)}{2})d]$. Thus the number of distinct sums a k element subset can have is $\frac{(n-k-1)(n-k)}{2} - \frac{(k-1)k}{2} + 1 = kn - k^2 + 1$. To get the total number of unique sums we carry out the summation as k varies from 1 to n to get

$$\sum_{k=1}^n (kn - k^2 + 1) = \frac{n^2(n+1)}{2} - \frac{n(n+1)(2n+1)}{6} + n = \frac{n^3 + 5n}{6}.$$

Now including the empty set sum of 0, we have $\frac{n^3+5n}{6} + 1 = \frac{n^3+5n+6}{6}$ distinct sums as claimed. $\square$

From Lemma 11.4, it follows that

$$\begin{aligned} |V_m| &\leq \frac{2U}{n(n+1)} |V_0| \\ &\leq \frac{n^3 + 5n + 6}{3n^2 + 3n} |V_0| \approx \frac{n}{3} |V_0|. \end{aligned}$$

Again the peak growth factor is significantly smaller than the worst case complexity.

The above two cases can be generalized to cases such as 2-CP (disjoint union of two CPs where $U = O(n^2)$) and 2-AP (disjoint union of two APs where $U = O(n^6)$). As U is made to increase by reducing the additive structure (by the union of a larger number of CPs or APs), will the complexity continue to increase? No, since we showed the maximum growth factor $\eta_{peak} \leq n^7/9$. So, for all instances with $U > O(n^7)$, we can expect no further increase.

11.3 A Special Case: $\mathbf{a} = \mathbf{b}$

Next, let's consider the special case of a Geometric Progression where $\mathbf{a} = \mathbf{b}$. In this case, there is no additive structure at all, since the elements are related by geometric growth. We have $U = 2^n$ and $N_{LO} = 1$. The loose bound of Lemma 11.4 is not useful in this case as it suggests that this could be a worst case. However, it turns out that this special case is trivially easy with regard to our algorithm. In this case, first note that all the NDPs fall onto a single straight line connecting the points $(0, 0)$ and (B_n, B_n) . See the bottom plot of Figure 18. Consider some edge $e_i \in q_k$ through which the OL passes through. This edge may interact with several edges of P_{k-1} in the next level. However, only one of these nodes will have a valid path. All others will have invalid paths, i.e., paths outside the bounds of the corresponding edges. This is so, since $\mathbf{a} = \mathbf{b}$ is *super-increasing* ($a_i > \sum_{j=1}^{i-1} a_j$, for all $i > 1$), leading to the non-overlap of subspaces S_j at $(0, 0)$ and S_j at (b_{j+1}, a_{j+1}) for any $j \in [n-1]$. As a result, the FILTER step prior to the iterations in Algorithm 5, will remove all paths except the solution path. The final graph is just a single path consisting of a sequence of at most n nodes, i.e., $|V_m| \leq n$.

11.4 A snapshot of instances

A rough understanding of the *degree of additive structure* can be obtained visually by looking at the geometry of the instance. Of course this is useful only for small instances. When the degree of additive structure is high the geometric structure is wide and short, and when the additive structure is low it is thin and long. Below we show the geometric structures of four instances (a *CP*, an *AP*, a *GP* and a random case) to provide a visual understanding at a glance in Figure 18. Comparing the random case to that of *CP*, we notice a *shearing* of the geometric structure. This *shearing* provides a rough qualitative reason on why the distinct local intersections are polynomially bounded; if we view the entire point set as a collection of smaller subspaces, for a given OL, many subspaces go out of range due to shearing. Thus the number of distinct intersections is smaller.




Figure 18: The geometric structures for four instances with varying degrees of additive structure. For visual clarity we used small sized instances ($n = 9$). Notice the shearing of the geometric structure in the *random* case compared to a *CP*.

12 Epilogue

In this section we make some philosophical observations on the main topic of the paper, and offer some personal opinions. It should be emphasized that these are not precise statements, but at best some rambling personal thoughts. These are divided into two parts: in the first part we discuss some reasons that *likely influenced* the majority opinion on the $\mathcal{P}$ vs. $\mathcal{NP}$ question, and in the second part the historical antecedents of the problem are discussed.

12.1 A combinatorial facade and a reconciliation

In many mathematical problems and particularly in combinatorics, the power set of some relevant base set is the universe. Whenever we think of all possible combinations (of elements of a set), we are implicitly thinking of them as *equal* in some sense. Each is just a combination out of the many possibilities, and contributing an unit to the total count. We are often required to estimate the number of combinations (satisfying some property) associated with a problem, but less concerned with how these are related to each other in the first place. This is particularly true when faced with a tremendously large number of combinations. In other words, we are intuitively coaxed into believing that either *no structure is possible*, or it is *futile to look for a structure*, but be content with the count. Thus, the *combinatorial explosion* eclipses any attempts to understand the structure.

For any $\mathcal{NP}$ -complete problem the *combinatorial aspect* presents itself readily by its hypercube with an exponential number of vertices. From the phrasing of the *Subset Sum Problem*, one sees it as an *one-dimensional problem* with the 2^n scalar numbers corresponding to the subset sums. For any given instance, at first glance we can rule out only a small number of vertices from the search space; after that almost all of the hypercube still remains. Furthermore, the hypercube doesn't reveal why one subset is more *likely* than another to be the solution. All of them appear equally likely and one is led to believe that the structure is *indeed* that of a combinatorial type. Perhaps these are some reasons why so many researchers intuitively felt that an exhaustive search cannot be avoided.

In our quest to understand the structure, we converted the *Subset Sum Problem* into a *two-dimensional problem* by attaching a unique index to each element of the problem. With this the hypercube is transformed into a point set in the Cartesian plane, and suddenly we see some structure that was not evident in the *one-dimensional* view³. Further efforts to fully understand the point set enabled us to characterize the structure of the *Subset Sum Problem*. We were able to find enough structure so that we can always solve the decision problem. Once we understood the structure we realized that it is the much smaller *relational aspect* that is important, and it was being shielded all along by a seemingly impregnable combinatorial facade.

³Regarding the power of an extra dimension, one can recall from Calculus how several integrals of a real variable become amenable to exact evaluation *only* when transferred to the complex plane.

12.2 Historical significance

The *Theory of Computing* is a vast subject spanning centuries, with many crises and triumphs, and featuring contributions from some of the greatest minds. There is no way to trace a short path through history without missing out on many important developments that eventually paved the way to the $\mathcal{P}$ vs. $\mathcal{NP}$ question. In this short section we only sample some key moments through history to provide an overly simplistic and rough historical perspective in a highly informal way.

The great mathematician David Hilbert, in 1928 [14], asked the question: “Is there a finite procedure to determine the truth or falsity of any mathematical statement”? What Hilbert had in mind was a “mechanical process” of applying inference rules to a set of axioms and checking the truth of a given statement. In modern terminology, his question is about the existence of an *algorithm* to *decide* the truth of any mathematical statement in a given axiomatic system. This famous question is known as the *Entscheidungsproblem* (decision problem). Hilbert believed that every mathematical question must have a definite answer, either by a solution or by a proof of impossibility. His guiding view was: “*Wir müssen wissen, wir werden wissen*” (We must know, we will know), that there can be no ignorance in mathematics, and that we can know everything.

It turns out that Hilbert was not the first to have such a grand ambition. More than 250 years prior, the great mathematician and philosopher Gottfried Wilhelm Leibniz, dreamt of mechanization of reason, which he called a “wonderful idea” (see Davis [9]). According to Davis [9], Leibniz wanted to collect all human knowledge, extract key concepts and attach a distinct character(symbol) to each. Then, using rules of deductive logic one could mechanically manipulate the symbols to arrive at truths. This extraordinarily ambitious project, could not be realized particularly given the times of Leibniz. But evidently Leibniz continued thinking and writing about it till the end.

However, the great logician Kurt Gödel in 1931 showed by his famous *incompleteness theorems* that in any consistent axiomatic system, there exist truths that are not provable (see for e.g the translation [12]). Given an axiomatic system, let J_1 be the set of all statements that are provable, and let J_2 be the set of statements that are not provable. Then Gödel’s theorems imply $J_2 \neq \emptyset$, but do not preclude the existence of a classification procedure, which given a statement s can tell if $s \in J_1$ or $s \in J_2$. Turing in 1936 [29], succeeded in formalizing the “mechanical process” of Hilbert, by defining *computability* using *Turing machines*. Our current digital world has been made possible by this singularly powerful result of Turing! Turing showed the equivalence of *computability* and *halting* of *Turing Machines*(*TMs*). Turing’s result says that the set J_1 corresponds to the set of *TMs* that *halt* in a finite amount of time. In a way, the same sets are relabeled into $J_1 \cong \text{HALT}$, and $J_2 \cong \text{DoNotHALT}$. He also showed that given a statement, whether the corresponding *TM* halts or not, is *undecidable*. In other words, given a statement, there is no procedure that can tell us if it is provable or not. These results effectively put an end to Hilbert’s dream of placing all of mathematics in a formal axiomatic system.

Out of the problems that are decidable, we are interested in those problems where a

claimed solution can be verified in polynomial time, which is the class $\mathcal{NP}$. Cook [6] discovered the first $\mathcal{NP}$ -complete problem in *Boolean Satisfiability Problem (SAT)* which has the property that every problem in $\mathcal{NP}$ can be reduced to it in polynomial time, and which is at least as hard as any problem in $\mathcal{NP}$. The class $\mathcal{NP}$, as it turns out is extremely rich in problems of interest to humanity covering diverse areas such as circuit design, mathematics, computer science, industry, artificial intelligence, cryptography, biology, economics, medicine, physics and more [31; 21]. Cook’s amazing discovery reduced thinking about a multitude of problems in $\mathcal{NP}$ to a single $\mathcal{NP}$ -complete problem! Thus, if we can efficiently solve *SAT* or equivalently any one of the thousands of other $\mathcal{NP}$ -complete problems, we can solve all problems in $\mathcal{NP}$ also efficiently. What is the worst case complexity of *SAT*? Thus, the $\mathcal{P}$ vs. $\mathcal{NP}$ problem was born.

All the $\mathcal{NP}$ -complete problems admit trivial exponential algorithms involving a brute force search over the exponential search space. As many efforts over several decades to find efficient algorithms to $\mathcal{NP}$ -complete problems weren’t successful researchers started to believe that exhaustive search cannot be avoided. Amazingly enough, Kurt Gödel anticipated the $\mathcal{P}$ vs. $\mathcal{NP}$ problem 15 years earlier in 1956 even before these classes were formally defined. He viewed it as a bounded version of Hilbert’s *Entscheidungsproblem*[1]. In his letter to John von Neumann, Gödel wrote: “Now, it seems to me, however, to be totally within the realm of possibility that $\varphi(n)$ grows slowly”[25]. Here $\varphi(n)$ refers to the number of steps taken by a Turing machine to decide the truth of a propositional formula limiting to proofs of length n . We can understand this to be the time complexity of a $\mathcal{NP}$ -complete problem in equivalent terms.

With our result that $\mathcal{P} = \mathcal{NP}$, certainly Gödel’s expectation has come true. However, Gödel also hoped for a small degree polynomial n^c where $c \leq 2$. We are not there yet, but with adequate improvements to the complexity, there is hope to realize the “wonderful idea” of Leibniz and fulfill Hilbert’s quest. We await to see what the future beholds! Going back to Hilbert’s optimism, our result serves as a vindication to a degree (after the negative results of Gödel and Turing), showing that he was correct after all as *we can know all the knowable truths*, at least those of the easily verifiable kind, in principle!

Dedication

To the memory of my loving father and high school mathematics teacher B. Venkateswarlu, who lit the initial flame, my first philosopher and personal hero, and to whom I will remain indebted forever; To the many great mathematicians of past and present who by their work, and the many selfless teachers who by their exposition, inspired in me a love for mathematics; To the Eternal One who inspires, enables and fulfills everything, I dedicate this work with a deep sense of gratitude. *Śrī Gurudeva-caraṇārpaṇamastu. Om Tat Sat.*

References

- [1] S. A. Aaronson. *P $\stackrel{?}{=}$ NP*. Springer International Publishing, 2016.
- [2] S. Arora and B. Barak. *Computational Complexity, A Modern Approach*. Cambridge University Press, 32 Avenue of Americas, New York, NY 10013-2473, 2009.
- [3] A. Becker, J-S. Coron, and A. Joux. Improved generic algorithms for hard knapsacks. *EUROCRYPT 2011*, 6632:364–385, 2011.
- [4] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal of Computing*, 26(5):1411–1473, 1993.
- [5] S. Cook. The P versus NP problem (Clay Math Institute official problem description). <http://www.claymath.org/sites/default/files/pvsnp.pdf>, 2000.
- [6] S. A. Cook. The complexity of theorem-proving procedures. *Proc. 3rd Ann. ACM Symp. on Theory of Computing*, pages 151–158, 1971.
- [7] J. M. Coster, A. Joux, A. B. Lamacchia, M. A. Odlyzko, C-P. Schnorr, and J. Stern. Improved low density subset sum algorithms. *Computational Complexity*, 2:111–128, 1992.
- [8] S. Dasgupta, C. H. Papadimitriou, and U. V. Vazirani. *Algorithms*. Addison-Wesley, Reading, Massachusetts, 2006.
- [9] M. Davis. *the Universal Computer*. CRC Press, Taylor and Francis Group, i6000 Broken Sound Parkway NW, Suite 300, Boca Raton, FL 33487-2742, 2018.
- [10] L. Fortnow. *The Golden Ticket: P, NP, and the Search for the Impossible*. Princeton University Press, 41 William Street, Princeton, New Jersey 08540, 2013.
- [11] R. M. Gary and S. D. Johnson. *Computers and Intractability*. W. H. Freeman and Company, New York, 1979.
- [12] K. Gödel. *On formally undecidable propositions of Principia mathematica and related systems (Translated by Meltzer B.)*. Dover Publications Inc., New York, (from Basic Books Inc., 1962), 1992.
- [13] J. Hartmanis. Gödel, von Neumann and the P=?NP problem. *Bulletin of the European Association for Theoretical Computer Science*, 38:101–107, 1989.
- [14] D. Hilbert and W. Ackermann. *Grundzüge der theoretischen Logik*. Springer Verlag, 1928.
- [15] E. Horowitz and S. Sahni. Computing partitions with applications to the knapsack problem. *Journal of the Association for Computing Machinery*, 21:277–292, 1974.
- [16] N. Howgrave-Graham and A. Joux. New generic algorithms for hard knapsacks. *EUROCRYPT 2010*, 6110:235–256, 2010.

- [17] M. R. Karp. Reducibility among combinatorial problems. *R. E. Miller and J. W. Thatcher (eds.), Complexity of Computer Computations*, pages 85–103, 1972.
- [18] J. C. Lagarias and A. M. Odlyzko. Solving low density subset sum problems. *Journal of the Association for Computing Machinery*, 32(1):229–246, 1985.
- [19] R. J. Lipton. *The P=NP Question and Gödel’s Lost Letter*. Springer Science and Business Media, 233 Spring Street, New York, NY 10013, 2010.
- [20] R. J. Lipton and Regan K. W. *People, Problems, and Proofs*. Springer-Verlag Berlin Heidelberg 2013, 2013.
- [21] C. Moore and S. Mertens. *The Nature of Computation*. Oxford University Press, Great Clarendon Street, Oxford OX2 6DP, 2015.
- [22] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, ISBN, 0201530821, 9780201530827, 1994.
- [23] R. Schroepel and A. Shamir. A algorithm for certain np-complete problems. *SIAM Journal on Computing*, 10(3):456–464, 1981.
- [24] A. Sipser, M. *Introduction to The Theory of Computation, Third Edition*. Cengage Learning, 20 Channel Center Street, Boston MA 02210, 2013.
- [25] M. Sipser. The history and status of the P versus NP question. *In Proc. 24th Symp. on Theory of Computing*, pages 603–618, 1992.
- [26] L. J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1977.
- [27] T. Tao and V. Vu. *Additive Combinatorics*. Cambridge University Press, 2006.
- [28] B. A. Trakhtenbrot. A survey of russian approaches to perebor (brute-force search) algorithms. *Annals of the history of Computing*, 6(4):384–400, 1984.
- [29] A. M. Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(236):230–265, 1936.
- [30] L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(4):189–201, 1979.
- [31] A. Widgerson. *Mathematics and Computation, A Theory Revolutionizing Technology and Science*. Princeton University Press, 41 William Street, Princeton, New Jersey 08540, 2019.

13 Appendix

In this section we provide some numerical results to illustrate the behavior of the algorithm and highlight key metrics for which upper bounds have been derived in the main paper. As pointed out at the beginning of the paper the algorithm has a fairly highly complexity. As a result of this, we could not simulate instances with large n due to high memory demands. Furthermore, we have not made adequate efforts to optimize the code. Nevertheless, the claims of the paper can be understood even with smaller instances.

13.1 Example 1: Multiple Solutions

We first present an example with $n = 40$ and $m = 23$, where we have multiple solutions. The sequence $\mathbf{a}$ is

(102347	237178	278185	428923	492462	508288	678198	706739
914229	921016	1069771	1117479	1135293	1161956	1320170	1349570
1563847	1645714	1710619	1726111	1778124	1997780	2129472	2132052
2351467	2537993	2613230	2878081	2943934	2992315	3096988	3580746
3633163	3758382	4064504	4274011	4638683	4914859	5518984	6399285),

the target number $T = 43,665,189$, and $A_n = 87,302,148$. For this instance of SSP, the initial orbital graph has $|V_0| = 140,090$ nodes and $|E_0| = 587,913$ arcs. Recall that we have the progression of graphs $G_0, G_1, \dots, G_m$ for the m iterations of the algorithm. The number of nodes and arcs as a function of iterations is shown in the top plot of Figure 19. The final graph G_m has 283,577 nodes and 316,691 arcs, and thus about 2 times larger compared to G_0 . The bottom plot shows the overall growth factor of nodes and arcs as a function of iteration. The *first* iteration at which we have the maximum number of nodes in the graph is denoted k_{peak} . In this example $k_{peak} = 9$ which is marked by a vertical dotted line. At iteration k_{peak} we have $|V_9| = 3,650,487$ and $|E_9| = 6,704,902$, and the maximum growth factor for nodes is about 26. The region to the left of the dotted line is denoted $R1$ where the graph size is increasing compared to the previous iteration. In region $R1$, more nodes are created (by REFINE) than removed (by FILTER) causing an increase in the graph size. However, the efficiency of filtering increases with each iteration, and at iteration k_{peak} the filtering step removes as many nodes as generated by refining. The region to the right of the dotted line is denoted $R2$ where the graph is non-increasing until only the zero paths remain. In this region, the filtering step removes at least as many nodes as those created by refining.

Some remarks follow:

- (i) The plots of Figure 19 are typical for any instance of SSP, and not for this example only. However, the location of $k_{peak} \in [m]$ and maximum growth factor depend on the instance and degree of additive structure.
- (ii) In this example, the final graph G_m has a size larger than the initial graph. This is typical when there are many solutions. It turns out that this example has 47,187




Figure 19: The number of nodes and arcs as a function of iteration is shown in the top plot. The bottom plot shows the growth factors $\frac{|V_k|}{|V_0|}$ and $\frac{|E_k|}{|E_0|}$ as a function of iteration k .

solutions (zero paths). If a given instance has just a few solutions, the region $R2$ is lowered significantly and gets closer to the x -axis. In the case where there are no solutions (the final graph $G_m = \emptyset$) it coincides with the x -axis. When there are large number of solutions, as in the case when m is small, the region $R2$ rises and becomes level with the peak graph size.

- (iii) Note that the peak growth factor $\eta_{peak} \leq 2^{k_{peak}}$. In this example $\eta_{peak} = 26.058$ which is much smaller than $2^{k_{peak}} = 512$. Also this much smaller than the theoretical maximum growth factor given by $\lceil 7 \log 40 \rceil = 38$.
- (iv) The plots show that an *energy barrier* must be crossed in order to discover the zero paths. This is an intrinsic and typical behaviour of our algorithm.

13.2 Example 2: A Single Solution

Next, we consider an instance with $n = 40$, $m = 50$ that has only one solution. The target value for this instance is $T = 11,734,810,597,199,265$, and the sequence $\mathbf{a}$ is

given below:

a = (5340838300510	72879613659014	106133278417124	188686744784933
196231635078036	204742082291972	213403962557486	244279857200551
255857319212728	273576871530631	289708754284101	334920772207496
372914717341816	429101169184127	441841927212719	450764719411565
461837234112443	496024035526859	514816372113660	537046917450189
562997054764154	640773760680379	643885354170988	705244168633211
707689919002051	787375951122724	806333455912404	815636078884372
854263299158483	878538656073276	883649606688863	915253093666743
917298576163677	964114532211852	1008165187539165	1014126287123623
1015148570166785	1033545047017119	1103557815472065	1123705011761031).

Recall that the x -coordinate of a *zero path* gives the *index* of the subset that sums to the target. In this example, the solution $index = 251872521694$, whose binary representation shows the elements:

$$[index]_2 = (0011101010100100110001011010110111011110).$$




Figure 20: The graph size and growth factor as a function of iteration for an instance with only one solution.

The various computed quantities are as follows:

$$\begin{aligned}
|V_0| &= 122695, & |E_0| &= 494892, \\
|V_{10}| &= 5785430, & |E_{10}| &= 10132432 \\
|V_m| &= |V_{13}| = 47, & |E_m| &= |E_{13}| = 46 \\
k_{peak} &= 10, & \text{and } \eta_{peak} &= 47.153 < 2^{10} = 1024.
\end{aligned}$$

Again the peak growth is significantly below the theoretical estimates, and the plots are shown in Figure 20.

13.3 Example 3: Growth of k_{peak} with n

As noted from the previous example, k_{peak} is a reasonable indicator of the maximum growth factor in the algorithm. In Figure 21 we show the variation of k_{peak} as a function of n . For each n we varied $m \in [10, 2n]$, and for each (n, m) pair we generated a random instance where each $a_i, i \in [n]$ was chosen randomly from the interval $[1, 2^m]$. The target value was chosen so that $T \approx \frac{A_n}{2}$ since the geometrical structure would have the widest span, in general, at half the total sum. This also corresponds to maximum possible orbital graph size. We took the $\max\{k_{peak}\}$ for each n and plotted it. We notice that it is a slow growing function much like a $c \log n$. For reference, we plotted $\lceil 7 \log n \rceil$ which is the theoretical bound we obtained for maximum number of distinct zero paths through any node. The plot of k_{peak} for these smaller values of n is much closer to $\lceil 3 \log n \rceil$ than the theoretical bound.




Figure 21: The variation of k_{peak} with n is shown by the blue curve. The theoretical bound of $\lceil 7 \log n \rceil$ is shown at the top for reference. The curve $\lceil 3 \log n \rceil$ is also shown for reference which is closer to data at these small values of n .

13.4 Example 4: Varying m

Next we consider multiple instances of size $n = 50$, but varying m . In each case the instance was generated by randomly picking numbers in $[1, 2^m]$. We plot the growth

curves for the various cases in the same plot to show the effect of m in Figure 22. Some remarks follow:

- (i) Recall that the peak growth factor corresponds to the maximum number of distinct zero paths through an edge. When m is very small, the peak growth factor is limited. As m increases, the growth factor increases as the edges can have more number of distinct zero paths.
- (ii) As m continues to increase beyond a threshold, the number of solutions decrease. Some of the paths will be non-zero paths which will be filtered away. Thus the curve drops and flattens out to support the number of zero paths for the given instance. In this plot $k_{peak} \approx 18$ which is less than the theoretical bound $\lceil 7 \log 50 \rceil = 40$.




Figure 22: Growth curves with varying m for $n = 50$.

13.5 Example 5: High degree of additive structure

In this example we consider the case $\mathbf{a} = AP(90, c, d)$ with $n = 90$, $m = 89$, and where

$$\begin{aligned} c &= 19,329,079,171,151,820,605,874,590, \\ d &= 5,835,760,507,709,645,161,289,732 \\ \text{and } T &= 12,559,207,227,207,892,175,122,963,296. \end{aligned}$$

The graph size and the growth factor are shown in Figure 23. In spite of the larger size of the problem, the peak growth factor is small due to the high degree of additive structure. The number of solutions is 120,723,382,126,197,997,657,472.




Figure 23: Progression of graph size and growth factor for an AP, where $k_{peak} = 3$.

13.6 Some Combinatorial Applications

The algorithm's ability to determine the number of solutions is very valuable. We can use it to determine many combinatorial quantities. Two examples are shown below. These examples also serve as additional proof of the correctness of the algorithm.

13.6.1 Binomial Coefficients

While there are well known efficient methods to computing Binomial coefficients, we can also use the SSP algorithm. To compute $\binom{n}{k}$, simply choose $\mathbf{a} = \mathbf{1}_n$ and $T = k$. As an example, for $n = 80$ and $k = 40$, using the SSP algorithm, we obtained the number of zero paths to be 107,507,208,733,336,176,461,620, which is the correct result.

13.6.2 Restricted Partitions I

Consider the problem of finding the number of distinct partitions of a given positive integer N where each part is at most $K \in \mathbb{N}$. We denote the number of partitions by $R(N, K)$. This is a well known *restricted partitions* problem for which there is no known closed form expression. Such quantities can be calculated by recursion following a *divide and conquer* approach. We can determine $R(N, K)$ using the SSP, by identifying $\mathbf{a} = [K]$ and $T = N$. Considering the problem of determining $R(915, 60)$, we obtained the number of zero paths to be

$$R(915, 60) = 3,360,682,669,655,028.$$

One can verify the correctness of the result by other methods.

13.6.3 Restricted Partitions II

This is like the previous case but we are interested in decomposing a given positive integer N into sum of distinct cubes, where each part is less than K^3 for some $K \in \mathbb{N}$. We denote this quantity by $R_{cubes}(N, K)$. Considering the problem of determining $R_{cubes}(12345, 50)$ we obtain that there are only 7 distinct decompositions into cubes. The solution indices are 76790, 79382, 80038, 90506, 141210, 142491 and 527286, whose binary form will give the elements that sum to the target. The solutions are shown explicitly (in the same order) below:

$$\begin{aligned} 12345 &= 2^3 + 3^3 + 5^3 + 6^3 + 7^3 + 8^3 + 9^3 + 10^3 + 12^3 + 14^3 + 17^3 \\ 12345 &= 2^3 + 3^3 + 5^3 + 10^3 + 11^3 + 13^3 + 14^3 + 17^3 \\ 12345 &= 2^3 + 3^3 + 6^3 + 8^3 + 12^3 + 13^3 + 14^3 + 17^3 \\ 12345 &= 2^3 + 4^3 + 8^3 + 9^3 + 14^3 + 15^3 + 17^3 \\ 12345 &= 2^3 + 4^3 + 5^3 + 8^3 + 9^3 + 10^3 + 11^3 + 14^3 + 18^3 \\ 12345 &= 1^3 + 2^3 + 4^3 + 5^3 + 8^3 + 11^3 + 12^3 + 14^3 + 18^3 \\ 12345 &= 2^3 + 3^3 + 5^3 + 6^3 + 8^3 + 9^3 + 10^3 + 12^3 + 20^3. \end{aligned}$$

Looking just at the first two solutions, we see that

$$6^3 + 7^3 + 8^3 + 9^3 + 12^3 = 11^3 + 13^3.$$

It should be clear that innumerable such relations can be extracted, in general, by having access to all solutions, and that many other combinatorial applications are possible.

13.7 A remark on size

In both the above examples, the complexity of finding $R(N, K)$ using other approaches such as Dynamic Programming or Recursion, will depend on the size of N for a given K , and hence can increase without bound. However, in our algorithm the complexity is limited to a polynomial in K , since as we showed, *size* does not matter beyond a threshold!