The Calkin-Wilf Tree: Theme and Variations

A Thesis
Submitted to the Faculty of
WORCESTER POLYTECHNIC INSTITUTE
in partial fulfillment of the requirements
for the Degree of Master of Science in

Applied Mathematics

By:

Ben Gobler

April 2023

APPROVED:

____

Professor Brigitte Servatius, Advisor

Abstract

In a recent publication [8], Jack E. Graver describes a method for computing terms in the Calkin-Wilf sequence. First, we explore an original method which uses continued fractions to evaluate and locate terms in the Calkin-Wilf sequence. Then, we extend the Calkin-Wilf tree to include all of the rational numbers exactly once each. Another generalization of the tree characterizes the relationship between rational numbers and continued fractions with integer coefficients. From a shift in perspective, we study infinite continued fractions and irrational numbers, and their relationships with Calkin-Wilf paths. The highly regarded result in this section is an original explanation for why irrational square roots of positive rational numbers have periodic continued fractions with palindromic coefficients. Finally, we exhibit a matrix analogue of the Calkin-Wilf tree and use its properties to conclude which irrational numbers have periodic continued fractions.

Acknowledgements

The success of this thesis would not have been possible without the support of Professor Brigitte Servatius and Professor Herman Servatius. Together they guided me through the research and writing process, and gave me the freedom to explore and discover beautiful mathematics.

I am deeply grateful to Guillermo Nuñez Ponasso for our many conversations on the topics of this thesis. A generous and inspiring mentor, his ideas and suggestions have shaped a prominent amount of my research.

I would also like to thank the WPI Mathematical Sciences department for endorsing my work and awarding this body of research with the Provost's MQP Award for Mathematical Sciences. I am grateful to the Mathematical Sciences community for offering continual support, encouragement, and everyday kindness.

Contents

1 Listing the Rational Numbers

The rational numbers are countable, which means that there exists an infinitely long list containing each rational number at least once. It is tempting to think that an infinite list should be large enough to include any set of numbers at least once. But, to one's surprise, the famous diagonal argument raised by Cantor demonstrates that there exist sets, such as the real numbers, which cannot be listed.

To prove that the rationals are countable, it suffices to show that the positive rationals are countable. That is because a list which includes each positive rational at least once can be transformed into a list of negative rationals via negation, and then, together with $0/1$, the two lists can be interlaced to produce a list of all of the rational numbers.

Undergraduate students are offered a proof that the positive rationals are countable which begins as follows: construct a table whose rows and columns are indexed by the positive integers; the entry in row i and column j is the fraction j/i. By definition, every positive rational can be expressed as the ratio of two positive integers, and thus is guaranteed to appear in the table at least once. By rotating this table forty-five degrees and reading the entries left to right, one level at a time, we obtain the desired list.

	1	2	3	4	5	$\cdots$
1	$\frac{1}{1}$	$\frac{2}{1}$	$\frac{3}{1}$	$\frac{4}{1}$	$\frac{5}{1}$	$\cdots$
2	$\frac{1}{2}$	$\frac{2}{2}$	$\frac{3}{2}$	$\frac{4}{2}$	$\frac{5}{2}$	$\cdots$
3	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{3}{3}$	$\frac{4}{3}$	$\frac{5}{3}$	$\cdots$
4	$\frac{1}{4}$	$\frac{2}{4}$	$\frac{3}{4}$	$\frac{4}{4}$	$\frac{5}{4}$	$\cdots$
5	$\frac{1}{5}$	$\frac{2}{5}$	$\frac{3}{5}$	$\frac{4}{5}$	$\frac{5}{5}$	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

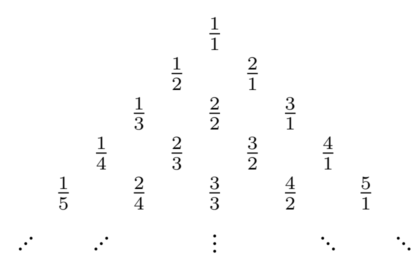

Each rational number appears infinitely often in the form of different fractions. Thankfully, this does not affect the countability argument; each number must appear in the list at least once, which is satisfied by infinitely many appearances. However, if we became interested in the more sophisticated task of listing each rational number *exactly once*, then we could choose to include only the fractions which are reduced. This modification is simple enough to implement: before adding a fraction to the list, check if it is reduced, and if not, discard it. This check can be performed by the Euclidean algorithm (see Section 4), since it is equivalent to show that the greatest common divisor of the fraction's numerator and denominator is 1. But if we want to find the trillionth term in this list, the suggested algorithm must write out the first trillion terms. In other words, we still lack a method for computing terms directly. In all cases, a direct algorithm requires the use of factorization, which is notoriously challenging (this is what makes encryption algorithms like RSA secure). One such implementation invokes the Euler totient function, whose formula uses the prime factors of a given integer. A full description of this method and others is given in [8].

We would prefer a list for which it is efficient to compute terms directly. In particular, we would like to answer two types of questions: a functional question such as, "What is the 200th number in the list?" and an inverse question such as, "Where does 22/7 appear in the list?" It is not clear whether any list can overcome the need for factorization in either algorithm. But it turns out that this is possible: we will describe a list and corresponding algorithms for both questions, each with logarithmic time complexity! Read on...

2 The Calkin-Wilf Tree

In their expository article [5], Neil Calkin and Herbert S. Wilf popularized what they called 'the tree of all fractions' (and others refer to as 'the Calkin-Wilf tree'). The same tree was considered three years prior [3], and other fraction-listing trees, such as the Stern-Brocot tree (see Section 9.1), had been known since the mid-nineteenth century. The Calkin-Wilf tree naturally solves the listing problem from Section 1 with the sought-after efficient algorithmic properties.

2.1 Construction

Start with a fraction p/q where p and q are positive integers. From this fraction, we produce two new fractions by either adding the numerator to the denominator or adding the denominator to the numerator. These operations are called the *left rule* and the *right rule*, respectively. When applied to the fraction 1/1, the left rule gives 1/2, and the right rule gives 2/1. We can apply the rules again to these fractions, and so on. The results of this recursive process are best organized in a binary tree, where the left child of p/q is $p/(p+q)$ and the right child is $(p+q)/q$, as given by the left and right rules. This is the *Calkin-Wilf tree*; its first four levels are shown below.

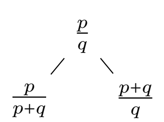

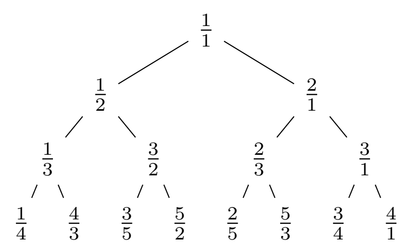

2.2 The Calkin-Wilf sequence

Our first observations of the Calkin-Wilf tree support a wonderful phenomenon: every fraction in the tree is reduced! Thinking ahead to potential listing algorithms, a method involving the Calkin-Wilf tree is already promising, since it entirely avoids the need to check for and delete fractions which are not reduced. And there is another astonishing property: every positive reduced fraction will appear somewhere in the tree, *exactly once*. That is, if you select your favorite positive rational, then it is guaranteed to appear uniquely in the tree, expressed as a reduced fraction. We will offer proofs of both properties in Section 5. Combined with the first property, we have that the Calkin-Wilf tree contains exactly the positive rational numbers, appearing once each.

By reading off the nodes of the tree, one level at a time, we obtain a list which includes each rational number exactly once. This list is called the *Calkin-Wilf sequence*, denoted $\ell(n)$, and it begins 1/1, 1/2, 2/1, 1/3, 3/2, 2/3, 3/1, 1/4, and so on. This is not the only way to arrange the nodes of a binary tree into a list, but it happens that the left-to-right approach takes advantage of a certain structural property of the binary tree which we will explore next.

The following progression of ideas was introduced by Graver in [8]. Consider a new tree derived from the Calkin-Wilf tree by replacing each fraction $\ell(n)$ with its term number, n. The root of the tree is 1, and its children are 2 and 3, and so on. Once we understand how to navigate this tree, we can locate term numbers and their corresponding fractions by traversing both trees simultaneously. What are the generating rules for the tree of term numbers? It can be shown by induction that the left and right children of n are $2n$ and $2n+1$, respectively, which one might agree is a delightfully simple pattern. In fact, the pattern is clearer when the term numbers are written in binary. Now, the left child of n is obtained by appending a '0' to n, and the right child is obtained by appending a '1'. That is because multiplication by 2 in binary is the bit shift operation. Thus, while the binary tree gets its name from the number of children at each node, it is a natural consequence that the same tree organizes the binary numbers so that traversing the tree corresponds with appending bits. Another interpretation is that each binary number in the tree encodes the series of directions from the root of the tree to its location, where a '0' indicates a left movement and a '1' indicates a right movement. The only subtlety is that the leading '1' does not represent any movement in the tree, since 1 is the root of the tree and is reached by no movements. For example, 1110 is reached from the root by two right movements and a left movement, corresponding with the three bits after the leading '1'.

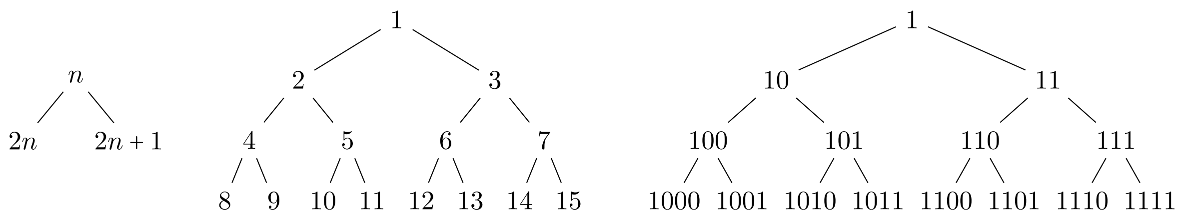

It seems inconsequential to ignore the leading '1's. They only cause a navigational hitch since the root of the tree is 1. To that end, might as well preprocess the binary numbers and remove their leading '1's; this would make the root of the tree the empty string, and every binary digit would correspond with a movement in the tree. But this rewriting is artificial and unnecessary; if the goal is to see that every binary digit corresponds with movement, then we should adjust the tree so that the root, 1, is a right child. Its parent could be the empty string, but this does not lend itself to arithmetic (what is 2 times the empty string?). A better choice is to label this node 0; although it breaks the pattern of appending (its right child is 1 instead of 01), it is mathematically motivated, since $2(0) + 1 = 1$. Likewise, we set $\ell(0) = 0/1$ because a right movement from 0/1 gives 1/1 by the right rule. This may appear to be more artificial than ignoring leading '1's, but it turns out that this choice dramatically simplifies the algorithm which computes $\ell(n)$ in Section 3.3.

With the trees extended by a node, we can use n to locate $\ell(n)$ in the Calkin-Wilf tree as follows: when written in binary, the digits of n indicate a sequence of left and right rules which, when applied to 0/1, evaluate to $\ell(n)$. For example, to find $\ell(14)$, write 14 in binary: 1110. Then on 0/1, perform the right rule three times, followed by the left rule once. This corresponds with the path $0/1 \to 1/1 \to 2/1 \to 3/1 \to 3/4$. So the 14th fraction in the list is 3/4, and we used 14 in binary as directions to navigate the Calkin-Wilf tree.

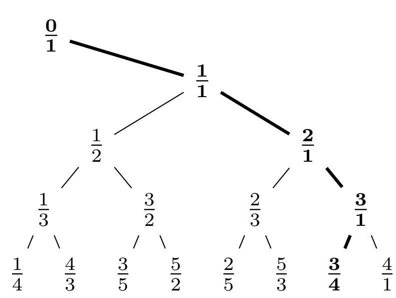

Note that this procedure also works for $n = 0$: applying the left rule to 0/1 gives 0/1, which correctly evaluates $\ell(0) = 0/1$. But this does not correspond with a path in the tree; there are no left movements from 0/1. This observation fits into a broader picture which we will develop in Section 6. For now, we can appreciate that our method for computing $\ell(n)$ works on all binary representations of n (that is, with any number of leading zeros) and does not have cases which require separate treatment.

3 Continued Fractions

3.1 Introduction

Continued fractions arise in many branches of mathematics, including number theory, geometry, and dynamical systems. Continued fractions first appeared as a result of the Euclidean algorithm (see Section 4), which has been known for thousands of years [10]. Colloquially, a continued fraction is a nested fraction with a particular structure, formed by addition and inversion in an alternating fashion. Here is an example of a continued fraction with value 67/52:

$$1 + \cfrac{1}{3 + \cfrac{1}{2 + \frac{1}{7}}}$$

It is easiest to evaluate a continued fraction like this one from the bottom up: first, add 2 and 1/7, then invert the result, and repeat. Notice that the numerators of this continued fraction are all '1's; this is a feature of *simple continued fractions*. From now on, for ease, we will write 'continued fraction' when we always mean 'simple continued fraction'. Since the numerators of a continued fraction are uniform, we only need to specify the sequence *coefficients*. It is conventional to write $67/52 = [1; 3, 2, 7]$, where the semicolon separates the first coefficient (representing the integral part of 67/52) from the remaining coefficients.

When the coefficients are restricted to positive integers (although the first coefficient may be zero), the continued fraction is called *regular*. Positive rational numbers have exactly two representations as regular continued fractions. In the example above, we can also write 67/52 as $[1; 3, 2, 6, 1]$. For reasons which will become clear in the next section, we always choose the representation which has an odd number of coefficients.

3.2 Continued Fractions in the Calkin-Wilf Tree

The narrative in Section 2 followed closely to the method developed by Graver in [8]. Now, we diverge with an original approach. Recall that we can compute $\ell(n)$ using a binary representation of n, where the '1's and '0's indicate right and left movements from 0/1 through the Calkin-Wilf tree. To improve efficiency, instead of performing the left and right rules individually, let's determine how to make consecutive movements simultaneously by augmenting the left and right rules.

We make m right movements from p/q by adding the denominator to the numerator m times. This gives the fraction $(p + mq)/q$, which simplifies to $m + p/q$. Likewise, we make m left movements from p/q by adding the numerator to the denominator m times, which gives $p/(mp + q)$. To simplify, we invert the reciprocal of this fraction, giving $1/((mp + q)/p)$. This is $1/(m + q/p)$, which, after one more double-inversion, is $1/(m + 1/(p/q))$. We call these modified rules the *m-right rule* and the *m-left rule*, respectively.

$$\frac{p}{q} \xrightarrow{m \text{ right}} m + \frac{p}{q} \qquad\qquad \frac{p}{q} \xrightarrow{m \text{ left}} \frac{1}{m + \frac{1}{\frac{p}{q}}}$$

Using these rules, we can traverse multiple levels of the Calkin-Wilf tree at once, as long as those movements are made in the same direction. All that remains is to combine these two types of movements. Consider a generic path through the Calkin-Wilf tree from 0/1 to p/q. We can use m-right and m-left rules to write p/q in terms of its ancestors at the turning points along the path. In the diagram below, these ancestors are labelled p'/q', then p''/q'', and so on.

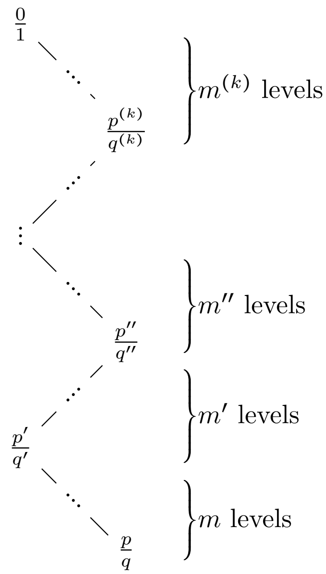

At the bottom of the diagram, p/q is m right movements from its ancestor p'/q'. From the m-right rule,

$$\frac{p}{q} = m + \frac{p'}{q'} \tag{1}$$

Likewise, p'/q' is m' left movements from its ancestor p''/q''. From the m-left rule,

$$\frac{p'}{q'} = \frac{1}{m' + \frac{1}{\frac{p''}{q''}}} \tag{2}$$

Combining (1) and (2) gives

$$\frac{p}{q} = m + \frac{1}{m' + \frac{1}{\frac{p''}{q''}}} \tag{3}$$

Now p/q is written in terms of p''/q''. We continue to iterate this process, writing p''/q'' in terms of its ancestors, and so on, until reaching an ancestor $p^{(k)}/q^{(k)}$ of the form $m^{(k)}/1$, which is $m^{(k)}$ right movements from 0/1. After this iteration, the expression for p/q is

$$\frac{p}{q} = m + \frac{1}{m' + \frac{1}{m'' + \frac{1}{\ddots + \frac{1}{m^{(k)} + \frac{0}{1}}}}} \tag{4}$$

And 0/1 is zero, so this simplifies to

$$\frac{p}{q} = m + \frac{1}{m' + \frac{1}{m'' + \frac{1}{\ddots + \frac{1}{m^{(k)}}}}} \tag{5}$$

The result is surprisingly simple: we can write p/q as a continued fraction whose coefficients are the numbers of consecutive left and right movements in the tree. This is why it is convenient for the paths to begin at 0/1; in the simplification from (4) to (5), the final ancestor 0/1 vanishes, and all that remains are the numbers of movements along the path. If instead we had chosen for the paths to begin at 1/1, then there would be cases for whether the path begins with a left or right movement, which makes the construction unnecessarily complicated. All this is streamlined by starting paths at 0/1.

Additionally, since $0/1$ is a right ancestor, $m^{(k)}$ counts right movements. But the construction begins by writing p/q in terms of its right ancestor, p'/q', so m also counts right movements. Therefore, the number of coefficients in the continued fraction produced by the construction is *always odd*. This is why in Section 3.1 we decided that of the two representations of p/q as a regular continued fraction, we will always choose the one with an odd number of coefficients because it is the one whose coefficients correspond with movements along the path in the Calkin-Wilf tree from $0/1$ to p/q.

3.3 Computing the Calkin-Wilf Sequence

To demonstrate the elegance of the method described in Section 3.2, let's compute $\ell(49)$. First, write 49 in binary: 110001. Now, instead of performing individual left and right movements for each of the bits, we partition the number into groups of consecutive '1's and '0's; this makes groups of sizes 2, 3, and 1. Thus, $\ell(49)$ is equal to the continued fraction with coefficients 1, 3, and 2:

$$\ell(49) = 1 + \cfrac{1}{3 + \frac{1}{2}} = \frac{9}{7}$$

We find that the 49th fraction in the Calkin-Wilf sequence is $9/7$.

There is one subtlety: the number of groupings may be even, in which case, the continued fraction produced by this method does not correspond with a path in the Calkin-Wilf tree. Since the first bit is a '1' (which corresponds with the initial right movement from $0/1$), an even number of groupings occurs when the final bit is a '0' (which corresponds with a final left movement). Certainly, paths in the tree can end with left movements. But the first coefficient in the continued fraction represents final right movements. Thus, in this case, we set the first coefficient m to be 0, indicating zero final right movements. This makes sense: the left rule creates proper fractions, so we should expect that a path ending with left movements leads to a fraction with integral part 0. For example, 24 in binary is 11000, which ends with zeros. Here is the result of the algorithm computing $\ell(24)$:

$$n = 24 \rightarrow 11000 \rightarrow \underbrace{11}_{2}\underbrace{000}_{3}\underbrace{}_{0} \rightarrow 0 + \cfrac{1}{3 + \frac{1}{2}} \rightarrow \ell(n) = \frac{2}{7}$$

The algorithm correctly obtains $\ell(24) = 2/7$.

Now, we make a crucial observation. Not only is this algorithm efficient, it is also reversible! That is, given a reduced fraction p/q, we can write it as a continued fraction with an odd number of coefficients (see Section 4.2) and then use the coefficients to write the term number of p/q in binary. Be cautious here: the first coefficient represents final movements, and the last coefficient represents initial movements, so the order of the coefficients appears reversed when written as the sizes of bit groupings. In the example below, the continued fraction coefficients 1, 3, and 2 become the groups sizes 2, 3, and 1.

$$n = 49 \leftrightarrow 110001 \leftrightarrow \underbrace{11}_{2}\underbrace{000}_{3}\underbrace{1}_{1} \leftrightarrow 1 + \cfrac{1}{3 + \frac{1}{2}} \leftrightarrow \ell(n) = \frac{9}{7}$$

The time complexity of this algorithm for computing $\ell(n)$ is linear in the number of bit groupings, which is at most the number of bits. Thus, the time complexity is $\mathcal{O}(\log n)$. To analyze the reverse algorithm, we need to know the complexity of writing a fraction as a continued fraction, which we'll answer in Section 4.2.

4 The Euclidean Algorithm

The Euclidean algorithm is a foundational tool in mathematics which computes the greatest common divisor of two numbers. The algorithm was first described by Euclid in his *Elements* in 300 BC and is one of the oldest algorithms still in use [10].

The algorithm works as follows: given two numbers a and b, there is a unique way to express a as $q \cdot b + r$ for nonnegative integers q and $r < b$. The number q represents the quotient, and r the remainder, of a upon division by b. The process is repeated with the two numbers b and r, and so on, until the remainder term is zero. Once this happens, the final number in the position of b is the greatest common divisor of the original two numbers. When the greatest common divisor is 1, we say that a and b are *coprime*, having only the trivial common factor.

4.1 The Euclidean Algorithm in the Calkin-Wilf Tree

The Calkin-Wilf tree is closely related to the Euclidean algorithm. Consider the right half of the Calkin-Wilf tree in the diagram below. The fraction 9/7 appears in the lowest level. Let's observe what happens when the Euclidean algorithm is performed on 9 and 7: after three iterations, the value of r is 0. Notice that the values of q, 1, 3, and 2, are the numbers of left and right movements along the path in the tree, and the values of b and r are the numerators and denominators of the fractions at the turning points.

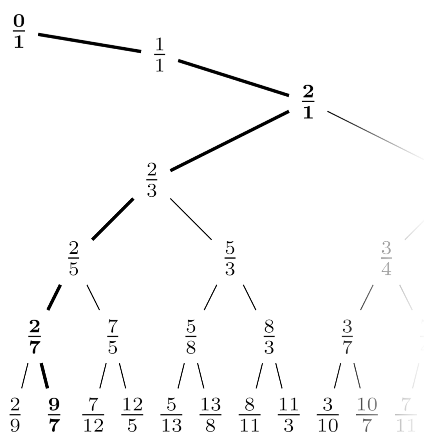

a	=	$q \cdot b$	+	r
9	=	$1 \cdot 7$	+	2
7	=	$3 \cdot 2$	+	1
2	=	$2 \cdot 1$	+	0

What is going on? To make the connection, it will help to understand the Euclidean algorithm in the same way that Euclid first described it. Rather than computing quotients and remainders via division or modular arithmetic, we can make the computation more elementary by performing the division with repeated subtraction. This is called the *slow Euclidean algorithm.* We will demonstrate that the path from 9/7 to 0/1 performs the slow Euclidean algorithm. At the bottom of the path, 2/7 is reached from 9/7 by subtracting 7 from 9. Since 2 is smaller than 7, the roles reverse, and the sequence $2/7 \to 2/5 \to 2/3 \to 2/1$ is obtained by repeatedly subtracting 2 from 7. Since 1 is smaller than 2, the roles reverse again, and 1 is repeatedly subtracted from 2, producing the sequence $2/1 \to 1/1 \to 0/1$. The algorithm terminates when this zero appears (another reason why it is useful to extend the tree by 0/1). And, that's it! In Section 5.1, we will explain why movements along an upward path correspond with subtractions. Finally, note that the number of repeated subtractions is exactly the quotient upon division, and what's left after the subtractions is exactly the remainder upon division. This explains why each quotient counts the number of consecutive movements in a particular direction, and the remainders appear in the fractions at the turning points.

Also, we see in the final line of this run of the algorithm ($2 = 2 \cdot \mathbf{1} + 0$) that the greatest common divisor of 9 and 7 is 1, which should be expected; every fraction in the tree is reduced, so the numerator and denominator pairs are all coprime.

4.2 Computing the Reverse Algorithm for the Calkin-Wilf Sequence

In Section 3.3, we discovered a reversible algorithm which can compute $\ell(n)$ given a term number n, and vice versa. The first step in the reverse algorithm requires us to write a given fraction p/q as a regular continued fraction with an odd number of coefficients. There is a simple way to do this which follows from the derivation in Section 3.2. Recall that the continued fraction coefficients of p/q are the numbers of left and right movements along the path from $0/1$ to p/q in the Calkin-Wilf tree. From Section 4.1, we know that the numbers of movements are exactly the quotient values when the Euclidean algorithm is run on p and q. So, the steps of the Euclidean algorithm produce the continued fraction coefficients of p/q. However, the Euclidean algorithm is not guaranteed to run for an odd number of steps. But that is fine; once the algorithm terminates, we can produce a continued fraction for p/q using the quotient values as coefficients, and then we can apply any necessary adjustments (as described in Section 3.1) so that the continued fraction has an odd number of coefficients.

In the example from Section 4.1, the run of the Euclidean algorithm on 9 and 7 reveals from the values of q that $9/7$ can be written as the continued fraction $[1; 3, 2]$, as we've seen before and can be quickly verified. However, the run of the algorithm on 7 and 2 only takes two steps (they are the final two steps from the run on 9 and 7). The result is $7/2 = [3; 2]$, which is accurate, but this continued fraction does not correspond with a path in the Calkin-Wilf tree since it has an even number of coefficients. Instead, we rewrite $[3; 2] = [3; 1, 1]$. Then the term number of $7/2$ in the Calkin-Wilf sequence is 10111, which is 23 in decimal. Thus, $\ell(23) = 7/2$.

Note that to run the Euclidean algorithm on p and q and expect a correspondence with p/q in Calkin-Wilf tree, it matters which order we input the numerator and denominator. In particular, we set $a = p$ and $b = q$. To see the distinction, let's run the algorithm with $2/7$, which we should expect to be different from the previous run on $7/2$. Note that $2/7$ is a proper fraction, so its continued fraction should have integral part 0. Can the Euclidean algorithm produce zeros? The answer is yes, and it happens naturally: 2 has zero multiples of 7, which is recorded by the algorithm. Then, in the next step, the positions of 2 and 7 are reversed! So, the algorithm is able to 'fix' the order of the numbers and run appropriately, recording a zero in the process which indicates that 2 is less than 7. In the end, we correctly obtain $2/7 = [0; 3, 2]$.

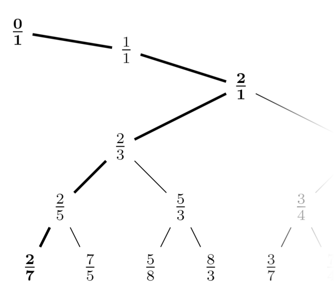

a	$=$	$q \cdot b$	$+$	r
2	$=$	$0 \cdot 7$	$+$	2
7	$=$	$3 \cdot 2$	$+$	1
2	$=$	$2 \cdot 1$	$+$	0

Finally, it is well known that the time complexity of the Euclidean algorithm when run on p and q is linear in the length of $p + q$, which is $\mathcal{O}(\log(p+q))$. This is left as an exercise and can be found easily by an online search. The worst runtime occurs when p and q are consecutive Fibonacci numbers, since each step of the algorithm makes only one movement in the tree; we will see these Fibonacci fractions again in Section 7.1. In summary, both the algorithm for computing $\ell(n)$ and the reverse algorithm have logarithmic time complexity, which is made possible by the use of continued fractions to solve the problem.

5 Variation 1: The Double Tree

We know how to navigate the Calkin-Wilf tree using left and right rules. But, after seeing the Euclidean algorithm run in Section 4, we may wonder how to algebraically travel back up the tree, undoing the left and right rules. Even earlier, we justified extending the tree by $0/1$ since it is a right rule away from $1/1$; can these types of extensions be made in general, and are they unique? We will answer these questions over the next two chapters.

5.1 Backward Movements

Our first observation is that the left and right rules are reversible: to undo a left movement, subtract the numerator from the denominator, and to undo a right movement, subtract the denominator from the numerator. Algebraically, the backward left rule sends p/q to $p/(q-p)$, and the backward right rule sends p/q to $(p-q)/q$. This is exactly how the slow Euclidean algorithm performed backward movements along paths the Calkin-Wilf tree in Section 4.1.

However, backward movements should be made with caution. Every node in a binary tree has two children, but, in contrast, each non-root node has only one parent. So, there is an asymmetry between constructing forward paths and backward paths. In particular, while there are many forward paths which begin at a fraction p/q, there is only one backward path which starts at p/q (we call this path the *ancestry* of p/q). Soon, we will use this fact to prove that the Calkin-Wilf tree contains each positive rational number exactly once. With that in mind, it is important to know how to tell if a given fraction is a left child or a right child. Once this is known, we can apply the appropriate backward rule to determine its parent. It turns out that no cleverness is needed: since the left rule produces proper fractions and the right rule produces improper fractions, we can immediately determine the identity of a given positive reduced fraction by comparing the sizes of its numerator and denominator (with the exception of $1/1$, since it, as the root of the Calkin-Wilf tree, has no parent).

For example, $9/7$ is an improper fraction, so it must be a right child. Its parent is obtained from the backward right rule by subtracting the denominator from the numerator, which gives $2/7$. As $2/7$ is a proper fraction, it is a left child, and, using the backward left rule, we find that its parent is $2/5$.

As an aside, we now have the necessary tools to prove both fundamental properties of the Calkin-Wilf tree that were claimed in Section 2.1. First, we will show that every fraction in the tree is reduced. This is straightforward with the following observation: if p/q is reduced, then so are $p/(p+q)$ and $(p+q)/q$. This argument can be made using modular arithmetic or contraposition. Therefore, if a fraction in the tree is reduced, so are its children. Inductively, it is enough to know that $1/1$ is reduced to conclude that every fraction in the tree is reduced. Next, to show that every positive reduced fraction appears exactly once in the tree, we will argue that every reduced fraction has a unique backward path to $1/1$ by directly constructing its ancestry. In the example above, we retraced two steps along the ancestry of $9/7$. Of course, there is nothing special about $9/7$; this can be done for any reduced fraction. The key to the argument is that the sum of the numerator and denominator strictly decreases after applying a backward rule. Since the fractions remain reduced at each stage (following the first proof, if a child is reduced, then its parent is reduced), and their numerators and denominators are positive, the process of recovering ancestry must eventually reach $1/1$, which has the smallest possible numerator and denominator sum. Thus, we have shown that each positive reduced fraction will appear in the Calkin-Wilf tree. And since the ancestry of p/q is forced at each step, the path from p/q to $1/1$ is unique, showing that each positive reduced fraction appears in the tree exactly once.

Now, let's think bigger. Both the Euclidean algorithm and our method for computing $\ell(n)$ involve paths in the tree which begin at 0/1. But nothing stops us from making more backward movements beyond 0/1. Let's try it! A backward right movement from 0/1 gives −1/1, and from there, more backward movements produce an entire Calkin-Wilf tree, upside-down, and all negative! We call this object the *double tree*.

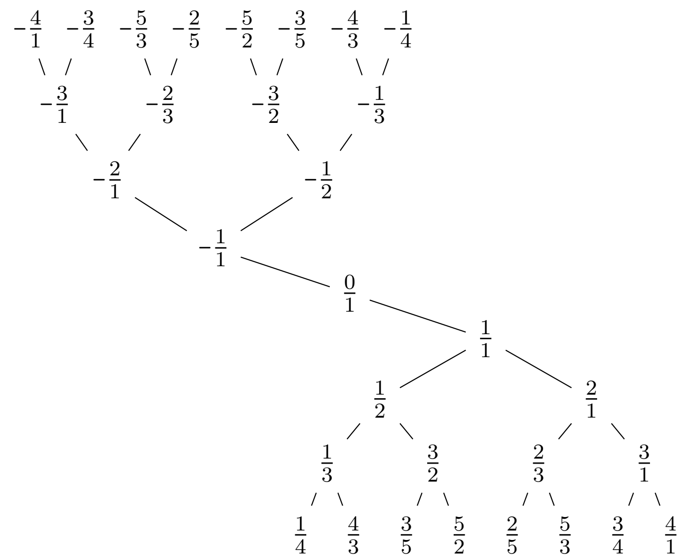

Although elaborate, the double tree is a much clearer picture for two reasons. First, it emphasizes that 0/1 is central to paths in the tree. We had this intuition when developing the correspondence between paths in the Calkin-Wilf tree and continued fractions (recall that the continued fraction coefficients of p/q correspond with the numbers of left and right movements along the path from 0/1 to p/q). Second, we can appreciate the appearance of the negative rational numbers. Not only do they also arrange into a tree, but it is essentially the *same tree* as the Calkin-Wilf tree, connected through 0/1. This has implications for our algorithm computing $\ell(n)$, which we will discuss next.

5.2 The Extended Calkin-Wilf Sequence

Until now, we chose to ignore the negative rationals after noting in Section 1 that a list of the positive rationals could be transformed into a list of all rational numbers in a straightforward way. In particular, given a list of the positive rationals, we can produce a list of the negative rationals by negating each term. Then, following 0, a list of the rationals is obtained by alternating terms from the positive and negative lists. While this new list certifies that the rationals are countable, it may have fewer structural properties from the perspective of a listing algorithm. For instance, we just secured an algorithm for computing $\ell(n)$ which elegantly converts term numbers to their corresponding fractions. If we use the interleaving technique to make a list of the rationals, then we will displace the positive rationals from their previous term numbers, rendering our algorithms useless. Ideally, we can avoid this by extending the list in a different way.

The Calkin-Wilf sequence is produced from the Calkin-Wilf tree by reading the nodes left to right, one level at a time. We can do the same for the double tree, except this produces a list which extends infinitely in two directions (sometimes called a 'doubly infinite sequence'). We call this the *extended Calkin-Wilf sequence*. Since it contains a copy of the Calkin-Wilf sequence, we choose to index the terms such that the positive rationals maintain their term numbers from the Calkin-Wilf sequence. Then 0/1 has term number 0, and −1/1 has term number −1, and so on.

$$
\begin{array}{ccccccccccccccccccccc}
n & = & \cdots & -8 & -7 & -6 & -5 & -4 & -3 & -2 & -1 & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & \cdots \\
\ell(n) & = & \cdots & -\frac{1}{4} & -\frac{3}{1} & -\frac{2}{3} & -\frac{3}{2} & -\frac{1}{3} & -\frac{2}{1} & -\frac{1}{2} & -\frac{1}{1} & \frac{0}{1} & \frac{1}{1} & \frac{1}{2} & \frac{2}{1} & \frac{1}{3} & \frac{3}{2} & \frac{2}{3} & \frac{3}{1} & \frac{1}{4} & \cdots
\end{array}
$$

The negative rationals have negative term numbers and are reached by backward movements from 0/1 in the double tree. We might wonder if the algorithm for computing $\ell(n)$ applies to these terms as well. The answer is yes, but with a few cosmetic changes.

Notice that, by symmetry, $\ell(-n) = -\ell(n)$. So, the algorithm could compute negative terms by just computing their positive counterparts and negating the results. This works, but there is a second interpretation which is worth understanding because it lends itself to the generalization in Section 6. Mirroring the derivation in Section 3.2, we would like to know how the continued fractions for negative rationals are related to their paths in the double tree. The path from 0/1 to, say, −9/7, consists of one backward right movement, then three backward left movements, and two backward right movements. Meanwhile, to write a continued fraction for −9/7, we can just negate the continued fraction for 9/7. But multiplying a continued fraction by −1 is equivalent to making each of the coefficients negative; this is done by distributing −1 to the denominator at each stage, as demonstrated below.

$$
-\left(a_0 + \cfrac{1}{a_1 + \cfrac{1}{\ddots + \frac{1}{a_k}}}\right) = -a_0 + \cfrac{1}{-\left(a_1 + \cfrac{1}{\ddots + \frac{1}{a_k}}\right)} = -a_0 + \cfrac{1}{-a_1 + \cfrac{1}{-\left(\ddots + \frac{1}{a_k}\right)}} = \cdots = -a_0 + \cfrac{1}{-a_1 + \cfrac{1}{\ddots + \frac{1}{-a_k}}}
$$

So from $9/7 = [1; 3, 2]$, we have $-9/7 = [-1; -3, -2]$. This indicates something remarkable: we can interpret backward movements as negative numbers of forward movements! That is, we might think to ourselves that the path to −9/7 consists of −1 right movements, then −3 left movements, and −2 right movements.

One last superficial change: our algorithm begins to compute $\ell(-n)$ by writing $-n$ in binary, but we have already trained ourselves to see the '1's and '0's as indicating forward movements. So instead, as if distributing the negation to each place value, we will make each bit negative. The '−1's and '−0's should be seen to indicate backward movements. Finally, the groupings of backward movements become the negative coefficients of the continued fraction for $\ell(-n)$. In the example below, the group sizes −2, −3, and −1 are the coefficients of the continued fraction for −9/7.

$$
\ell(-49) = -\ell(49) = -\frac{9}{7} = -\left(1 + \cfrac{1}{3 + \frac{1}{2}}\right) = -1 + \cfrac{1}{-3 + \frac{1}{-2}}
$$

$$
\ell(-49) = -1 + \cfrac{1}{-3 + \frac{1}{-2}} \;\leftrightarrow\; \underbrace{-1 \mid -1}_{-2} \mid \underbrace{-0 \mid -0 \mid -0}_{-3} \mid \underbrace{-1}_{-1} \;\leftrightarrow\; -110001 \;\leftrightarrow\; n = -49
$$

The only reason for this syntactic fussiness is to emphasize the connection between backward movements and negative continued fraction coefficients. In practice, none of these rewriting steps are necessary. The first method which we proposed for extending the algorithm computes negative terms by considering their positive counterparts and negating the results. Still, the important takeaway is that we have started to interpret backward movements as negative numbers of forward movements. We will make this more precise in Section 6.2. What appears to be a small observation about backward movements will generalize the relationship between continued fractions and paths in Calkin-Wilf-like trees.

6 Variation 2: The Four-Way Tree

In Section 5.1, we remarked that nodes in a binary tree have two children but only one parent. So, when constructing the ancestry of p/q in the Calkin-Wilf tree, we need to determine whether p/q is a left or right child in order to use the appropriate backward rule. But, thinking outside the box, there is nothing mathematically preventing us from applying any of the backward rules to any fractions in the tree. So, bravely, let's try it!

6.1 Construction

Starting at 0/1, we can make four types of movements: forward left, forward right, backward left, and backward right. Each movement produces a new fraction. Keep iterating, performing the four types of movements from these fractions, and so on. We call the resulting object the *four-way tree*.

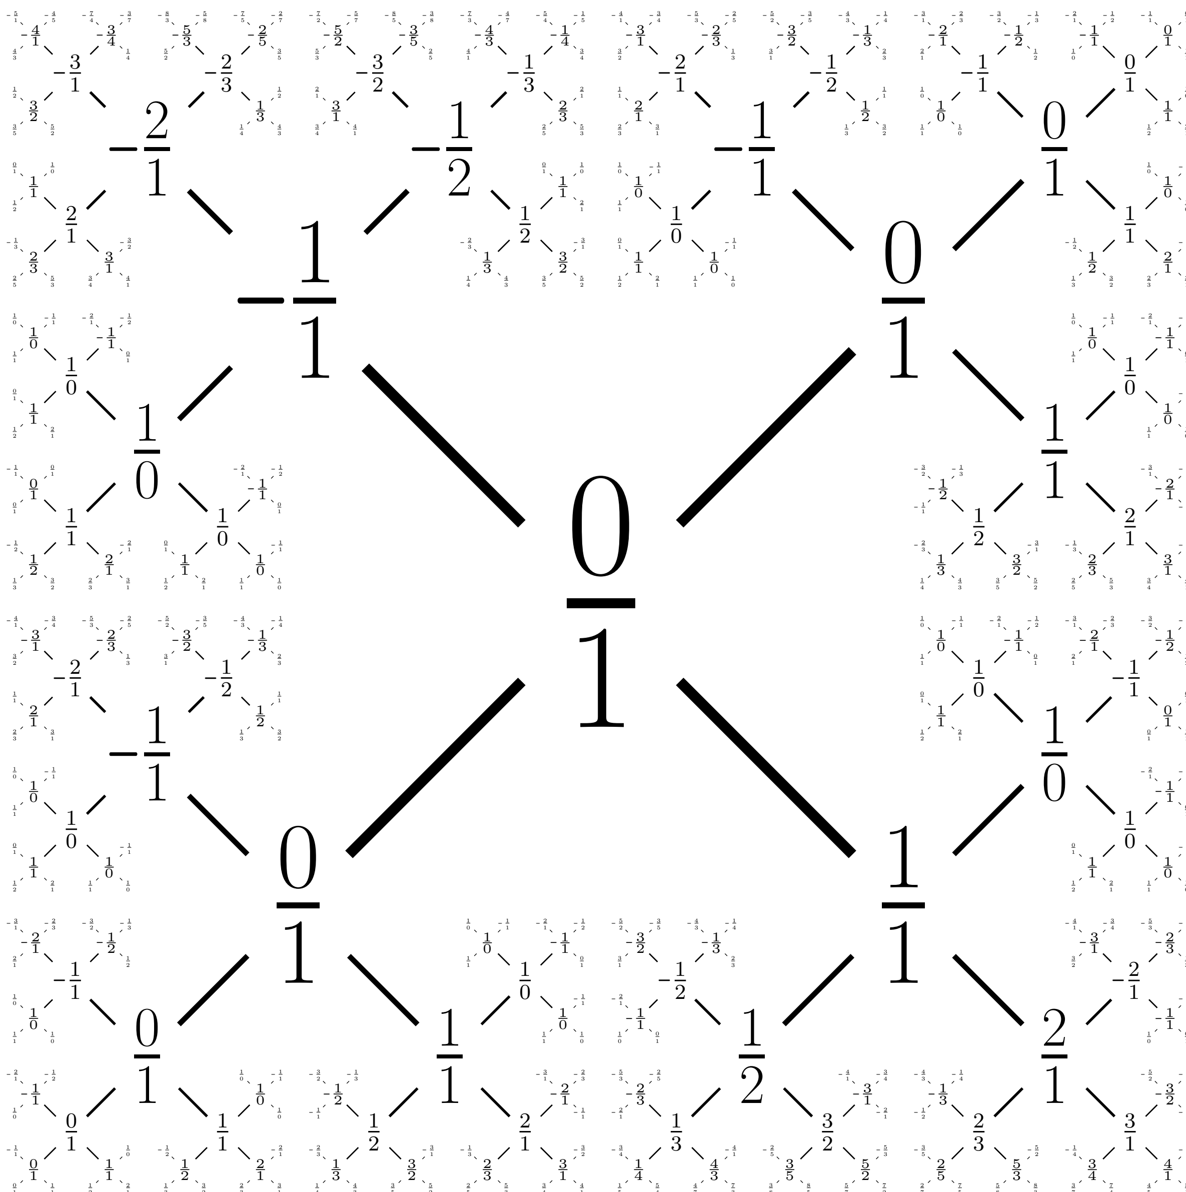

Wow! The four-way tree is a piece of art in itself. There are many patterns to observe, so let's start with what is most apparent. First, this tree contains repeated instances of the same fractions. In hindsight, this is necessarily the case; the double tree already contains each rational number, so any movement off of the double tree must lead to a repeated fraction. Prominently featured along the antidiagonal, we are reminded that the left rule sends 0/1 to itself. A careful eye may also spot appearances of the fraction 1/0: we will address this in Section 6.3.

We should be eager to know the significance of paths in this tree. For instance, all of the Calkin-Wilf and double tree paths embedded in the four-way tree correspond with continued fractions. We might wonder if all paths in the four-way tree have this property under some generalization of what we know about strict-forward and strict-backward paths. It turns out that, yes, all paths correspond with continued fractions. To understand how, we must pay closer attention to the relationship between forward and backward movements. While geometrically distinguished, we will find that they are formulaically identical.

6.2 Continued Fractions in the Four-Way Tree

When analyzing paths in the Calkin-Wilf tree, our first step toward the appearance of continued fractions was to algebraically construct the m-left and m-right rules, each capable of performing m consecutive movements in the same direction. Using this approach for the backward rules, we obtain the m-backward-left rule and the m-backward-right rule.

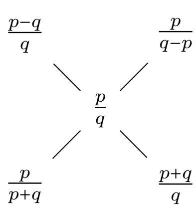

$$\frac{p}{q} \xrightarrow[]{\substack{\text{back}\\ m\text{ left}}} \frac{1}{-m+\frac{1}{\frac{p}{q}}} \qquad \frac{p}{q} \xrightarrow[]{\substack{\text{back}\\ m\text{ right}}} -m+\frac{p}{q}$$

$$\frac{p}{q} \xrightarrow{m\text{ left}} \frac{1}{m+\frac{1}{\frac{p}{q}}} \qquad \frac{p}{q} \xrightarrow{m\text{ right}} m+\frac{p}{q}$$

This looks promising: the only difference between the forward and backward rules is that, in both cases, m is replaced by $-m$. As suggested in Section 5.2, if we think of backward movements as negative numbers of forward movements, then the two types of rules are equivalent. For example, m-backward-left rule is the $(-m)$-left rule. It follows immediately that the results from Section 3.2 apply to all paths in the four-way tree. That is, a path from 0/1 to p/q corresponds with a continued fraction for p/q whose coefficients count the numbers of left and right movements along the path. In particular, forward movements correspond with positive coefficients, while backward movements correspond with negative coefficients. For example, here is a path in the four-way tree from 0/1 to 7/18 consisting of 3 right, 2 left, 2 backward right, 1 backward left, and 1 right movement. From this path, we write a continued fraction for 7/18 with corresponding coefficients.

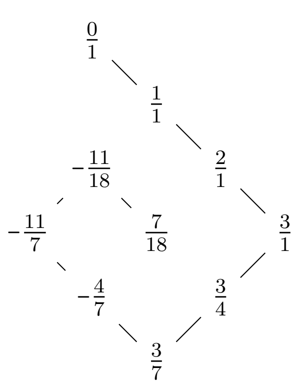

$$\frac{7}{18} = 1 + \frac{1}{-1+\frac{1}{-2+\frac{1}{2+\frac{1}{3}}}}$$

This is a pretty picture; each path in the four-way tree from $0/1$ to p/q corresponds with a continued fraction for p/q with integer coefficients. Conversely, if p/q can be written as such a continued fraction with an odd number of coefficients, then the coefficients describe a path in the tree from $0/1$ to p/q, where positive coefficients indicate forward movements, and negative coefficients indicate backward movements.

However, there are two subtleties to consider. First, unlike in the double tree, where the only initial right movements could be made from $0/1$, it is possible in the four-way tree to begin a path with left movements (albeit these movements are useless, since they take $0/1$ to itself). This is no trouble; as long as we remember that the final continued fraction coefficient represents initial right movements, we may use 0 to indicate that a path which begins with left movements has zero initial right movements. You may sense the danger of allowing 0 as a coefficient in the denominator of a continued fraction; we will address this in Section 6.3. The other subtlety is that continued fractions with integer coefficients are not always as well-behaved as regular continued fractions; we will explore this in Section 7.5.

6.3 Arithmetic with 1/0

The caveat to evaluating continued fractions with integer coefficients is that it is possible to arrive at division by zero. We also observed appearances of $1/0$ in the four-way tree. Instead of shying away from these occurrences of division by zero or making special cases for them, we can retain the result from Section 6.2 by specifying an appropriate convention for addition and inversion with $1/0$. Naturally, we define $m + 1/0$ to be $1/0$ for any integer m, and we define the inverse $1/(1/0)$ to be $0/1$, which is zero. Following these conventions, we can evaluate any continued fraction with integer coefficients, and remarkably, its value does appear at the end of the corresponding path in the four-way tree. In other words, the four-way tree supports these arithmetic rules. For example, the continued fraction below corresponds with a path from $0/1$ to $-2/1$ in the four-way tree. Accordingly, its rational value, found using these conventions, is -2.

$$-2+\cfrac{1}{3+\cfrac{1}{1+\cfrac{1}{-2+\frac{1}{1}}}} \;=\; -2+\cfrac{1}{3+\cfrac{1}{1+\frac{1}{-1}}} \;=\; -2+\cfrac{1}{3+\frac{1}{0}} \;=\; -2+\cfrac{1}{\frac{1}{0}} \;=\; -2+0 \;=\; -2$$

Here are some more examples involving $1/0$ which can be verified using the four-way tree in Section 6.1.

$$\frac{0}{1}=0+\cfrac{1}{-2+\frac{1}{0}} \qquad \frac{1}{1}=0+\cfrac{1}{-1+\cfrac{1}{1+\cfrac{1}{-1+\frac{1}{-1}}}} \qquad \frac{1}{0}=0+\cfrac{1}{-1+\cfrac{1}{1+\cfrac{1}{2+\frac{1}{0}}}}$$

We have assembled a complete narrative which relates rational numbers and (finite) continued fractions with integer coefficients. Every rational number has the same story: its reduced fractional form p/q appears exactly once in the double tree, and each path from $0/1$ to p/q in the four-way tree corresponds with a continued fraction for p/q with an odd number of integer coefficients. The only unusual detail is that we may, at times, require the use of $1/0$ and its associated arithmetic conventions.

But this is just scratching the surface; all of the techniques and results that have been developed so far will become useful guides when we extend our setup beyond the rational numbers. In fact, with a slight shift in perspective, we can use the similar ideas to draw conclusions about irrational numbers and more! The best results are yet to come. Stay tuned. . .

7 Variation 3: Exotic Calkin-Wilf Paths

So far, we have kept exclusively to the rational setting. The natural next step is to consider *irrational* numbers. While irrational numbers also have continued fractions and Calkin-Wilf-like trees, there are subtleties which require more careful attention. So, before making connections to the Calkin-Wilf tree, our first task is to become familiar with the relationship between irrational numbers and continued fractions.

7.1 Convergents

For clarity, let's start as we did in Section 1 by considering only the positive numbers. It is still the case that everything we discover about positive numbers applies equally to the negative numbers via negation; we will be more precise about this notion when it is relevant.

For show, here is the continued fraction for pi, an irrational number, which begins $[3; 7, 15, 1, 292, \dots]$.

$$\pi = 3 + \cfrac{1}{7 + \cfrac{1}{15 + \cfrac{1}{1 + \frac{1}{292 + \ddots}}}}$$

Unfortunately, our first observation appears to be an obstacle: since regular continued fractions for positive rational numbers are finite (think of the Calkin-Wilf tree), it follows that regular continued fractions for positive irrational numbers must be infinite. This is troubling from the perspective of computation. In particular, infinite continued fractions cannot be evaluated in the traditional sense (recall that continued fractions are evaluated from the bottom up) because they are bottomless. So, in order to assign values to these continued fractions, we'll need to establish a method other than direct computation.

This is similar to the issue of evaluating infinite series: one cannot sum infinitely many terms in a lifetime. Instead, we study the sequence of partial sums, obtained by truncating the sum after finitely many terms. If the sequence of partial sums converges, then we assign its limiting value to the infinite series.

We can extend this idea to infinite continued fractions by considering the sequence of *convergents*. This sequence is obtained by truncating the continued fraction after finitely many coefficients. Each term is a finite continued fraction, which can be evaluated directly. For example, consider the infinite continued fraction whose coefficients are all '1's: $[1; 1, 1, \dots]$. The first five convergents are shown below. In fact, this sequence converges to the golden ratio $\varphi = (1 + \sqrt{5})/2$ which, as expected, is irrational. So, we set $[1; 1, 1, \dots] = \varphi$.

$$1 \qquad 1 + \frac{1}{1} \qquad 1 + \cfrac{1}{1 + \frac{1}{1}} \qquad 1 + \cfrac{1}{1 + \cfrac{1}{1 + \frac{1}{1}}} \qquad 1 + \cfrac{1}{1 + \cfrac{1}{1 + \cfrac{1}{1 + \frac{1}{1}}}} \qquad \cdots \qquad 1 + \cfrac{1}{1 + \cfrac{1}{1 + \ddots}}$$

$$[1] \qquad [1; 1] \qquad [1; 1, 1] \qquad [1; 1, 1, 1] \qquad [1; 1, 1, 1, 1] \qquad \cdots \qquad [1; 1, 1, \dots]$$

$$1/1 \qquad 2/1 \qquad 3/2 \qquad 5/3 \qquad 8/5 \qquad \cdots \qquad \varphi$$

What if the convergents do not converge? For now, we are in luck: it is a foundational result in the theory of continued fractions that for any infinite regular continued fraction, its sequence of convergents always converges to an irrational number. Moreover, every irrational number has a unique infinite continued fraction (it no longer makes sense to talk about an even or odd number of coefficients). Proofs for both results can be found in [10]. This guarantee of convergence will cover us until Section 7.4, where we will eventually consider infinite continued fractions with integer coefficients and find them to be less well-behaved.

7.2 Periodic Paths

It is tempting to think about infinite regular continued fractions as corresponding with infinitely long paths in the Calkin-Wilf tree. But this interpretation is weak for two reasons. First, the Calkin-Wilf tree consists of arbitrarily long paths, but not infinitely long paths. After all, it does not make sense to follow a path forever, or to discover some number at the end of a path which never ends. Second, and more problematic, the coefficients of an infinite regular continued fraction do not describe the path in the tree that one might anticipate. For example, from the continued fraction of pi in Section 7.1, we might think that an infinite path to pi begins with 3 right movements, then 7 left movements, and so on. But this is wrong; remember that the first continued fraction coefficient represents the *final movements* along the path! So, we cannot even begin to draw the infinite path for a given infinite continued fraction because we do not know it starts, only how it ends (after an infinite sequence of earlier movements). Does all of this sound ridiculous? It should, and it felt preposterous to write. Again, infinite paths are an enticing first thought, but it is difficult to make the notion well-defined.

However, there is one idea regarding infinite paths which is barely more sensible: since the convergents of an infinite regular continued fraction are finite, each one corresponds with a finite path in the Calkin-Wilf tree. It may be that two convergents' paths have different initial segments. However, we might find a subsequence of paths which agree on initial movements. If such a subsequence exists, then its limit appears to be an infinite path which leads to the limiting value of the convergents. For concreteness, consider the golden ratio example from Section 7.1. Here are the paths from 0/1 to the first five convergents:

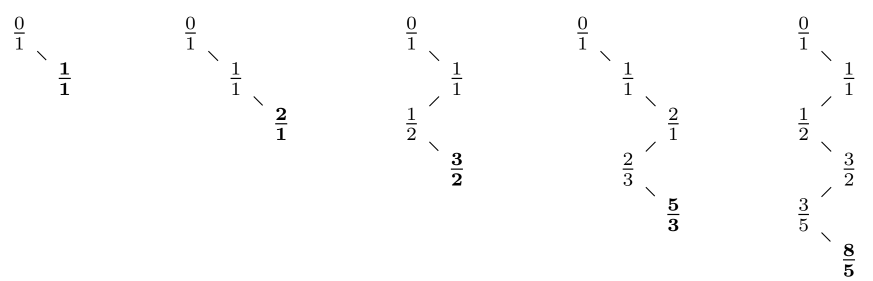

See that, separately, the odd-numbered and even-numbered convergents have paths which share initial segments, so they form two subsequences whose limits may be called infinite paths to φ. Whether or not we choose to adopt the notion of an infinite path, it is important to recognize that two convergents' paths can only share initial segments if the sequence of continued fraction coefficients repeats itself. This is certainly the case for φ, whose continued fraction coefficients are repeating '1's. More can be said about these types of paths and their initial segments, but we will refrain; it turns out that there is an even better model of this behavior, which we will explore next.

In Section 5.1, we proved that the Calkin-Wilf tree contains every positive rational number exactly once. So, if x is a positive rational number, then x appears in the Calkin-Wilf tree exactly once. Now, here's a twist: by contraposition, if x were such that it would appear more than once in the Calkin-Wilf tree, then x is not rational. That's an interesting piece of logic, but how would a number appear more than once in the tree? Well, we can construct such numbers directly.

Begin by choosing a finite segment of path consisting of left and right movements. Label the top of the path $x/1$. Then, as if this segment of path belongs to the Calkin-Wilf tree, use the left and right rules to produce fractions at the remaining nodes. Finally, we force x to repeat in the tree by setting the other end of the path equal to x again. Since x appears at both ends, we call the segment of path a *periodic path* for x.

The imposed equation can be solved for x to discover which number corresponds with this periodic path. For example, the golden ratio corresponds with the periodic path consisting of one left movement and one right movement. So, if the golden ratio appeared in the Calkin-Wilf tree, then it would repeat along this path. Therefore, as we already know, the golden ratio is irrational and does not appear in the Calkin-Wilf tree.

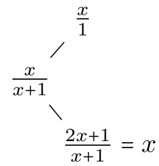

$$x^2 - x - 1 = 0$$
$$x = \varphi$$

But there's more: since this path was produced using left and right rules, we can apply our earlier results connecting Calkin-Wilf paths and continued fractions. Recall from Section 3.2 that we can write p/q as a continued fraction in terms of one of its ancestors and the number of movements along the path from p/q to that ancestor. In the special case where that ancestor is $0/1$, the continued fraction has only the numbers of movements (since $0/1$ is zero; this is why we extended the tree by $0/1$). In the periodic path above, we can write x in terms of its second ancestor, which is x itself, using a continued fraction. In particular, x is one final right movement away from its closest ancestor, which is one left movement away from its ancestor, x. Combining the m-right and m-left rules with $m = 1$ for both gives a continued fraction for x in terms of itself. We can recurse this definition to obtain an infinite continued fraction for x. This continued fraction is called *periodic* since its coefficients follow a repeating sequence. Notationally, we write $x = [\overline{1}]$, akin to the notation for repeating decimals. Moreover, we will call the path *regular* since it produces a regular continued fraction.

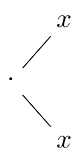

$$x = 1 + \cfrac{1}{1 + \frac{1}{x}} \qquad x = 1 + \cfrac{1}{1 + \cfrac{1}{1 + \frac{1}{1 + \ddots}}}$$

This continued fraction matches exactly with the result from Section 7.1. Hopefully, at first sight, you'll agree that this method is powerful; after choosing any segment of path, we obtain a corresponding irrational number, its minimal polynomial, and its regular continued fraction.

Here is another example of a regular periodic path. From $\frac{3x+4}{2x+3} = x$, we solve for x and discover that this is a periodic path for $\sqrt{2}$. The x at the bottom of the path is one right movement away from its ancestor, which is two left movements away from its ancestor, which is one right movement away from x again. Using this, we build a continued fraction for x in terms of itself, and the recursion produces an infinite continued fraction. Thus, $\sqrt{2} = [1; 2, 2, 2, \ldots] = [1; \overline{2}]$. Even though the first coefficient is not part of the period, this continued fraction is still called periodic, since the term refers to continued fractions whose coefficients are eventually repeating from some point onward. And, of course, $\sqrt{2}$ is irrational. One may appreciate this argument of irrationality in lieu of the traditional proof by contradiction.

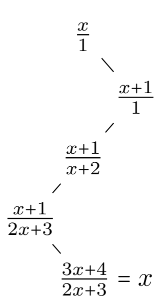

$$x^2 - 2 = 0$$
$$x = \sqrt{2}$$

$$x = 1 + \cfrac{1}{2 + \frac{1}{1+x}} \qquad x = 1 + \cfrac{1}{2 + \cfrac{1}{2 + \frac{1}{2 + \ddots}}}$$

Before moving on, we must address one technicality with this periodic path approach. A potential critique is the following: when we assume toward contradiction that x is a rational number (which is contradictory since x will be forced to appear twice in the Calkin-Wilf tree), shouldn't we write x as p/q, a ratio of whole numbers? Why are we allowed to write $x/1$ and perform the left and right rules on this fraction? The answer is simple, and it relies on a feature of the left and right rules which we have not yet taken advantage of.

Start with a fraction p/q, where p and q are integers. Now, multiply the numerator and denominator by a nonzero scalar m. What happens when we perform the left and right rules on $(mp)/(mq)$? We get $(mp)/(mp+mq)$ and $(mp+mq)/(mq)$, respectively. Factoring out m from the numerators and denominators, these are exactly $p/(p+q)$ and $(p+q)/q$, the left and right children of p/q. This may seem pointless, multiplying by m only to factor it out, but the success of the manipulation relies on the structure of the left and right rules. For example, if the left rule were to square the numerator, this argument would not hold. What we learn is that the left and right rules can be applied to $(mp)/(mq)$ for any nonzero m, and the results will be equal to the children of p/q. Setting $m = 1/q$ (where q is nonzero), we get $(p/q)/1$, which is the arrangement used when writing $x/1$ at the top of each periodic path.

With that settled, we should turn our attention back to the examples of periodic paths for φ and $\sqrt{2}$. Notice that in both examples, the imposed equation for x is a quadratic (this is not always the case; see Section 7.4). But quadratic equations have *two* solutions! Are the other solutions also represented by these periodic paths? In fact, they are. The key insight is that each periodic path leads from x to x in both directions. So far, we have only considered periodic paths as forward paths in the Calkin-Wilf tree. But since rational numbers appear exactly once in the double tree, our argument by contradiction applies to backward paths as well. That is, if x repeats along a backward path, then x is not rational. So, we missed half of the story! Let's go back for these second solutions.

The second solution to $x^2 - x - 1 = 0$ is $\phi = (1-\sqrt{5})/2$, which is the less famous cousin of the golden ratio. Now, let's retrace the ancestry of the upper x, which is a sequence of backward movements: there are no final right movements, but there is one backward left movement after one backward right movement from the lower x. Thus, the continued fraction coefficients are 0, −1, and −1. Recursing, we get the continued fraction $[0; -1, -1, -1, \dots] = [0; \overline{-1}]$ whose convergents converge to ϕ.

$$x \;\rightarrow\; x = 0 + \cfrac{1}{-1 + \frac{1}{-1+x}} \;\rightarrow\; x = 0 + \cfrac{1}{-1 + \cfrac{1}{-1 + \frac{1}{-1 + \ddots}}} = \phi$$

$$x \;\rightarrow\; x = 1 + \cfrac{1}{1 + \frac{1}{x}} \;\rightarrow\; x = 1 + \cfrac{1}{1 + \cfrac{1}{1 + \frac{1}{1 + \ddots}}} = \varphi$$

Finally, since the left and right rules can be applied along the periodic path in either the forward or backward directions, we might wonder why both computations impose the same quadratic equation for x. A nice explanation is that we can transform one of the continued fraction equations for x into the other. This is done by transferring a coefficient from one side to the other, then inverting both sides, and repeating. At certain steps in this process, the two sides of the equation represent the ancestry of either x to the same node along the path. The node begins as the lower x and transitions, one movement at a time, to the upper x.

$$x = 1 + \cfrac{1}{1 + \frac{1}{x}} \;\rightarrow\; -1 + x = \cfrac{1}{1 + \frac{1}{x}} \;\rightarrow\; \frac{1}{-1+x} = 1 + \frac{1}{x} \;\rightarrow\; -1 + \frac{1}{-1+x} = \frac{1}{x} \;\rightarrow\; \cfrac{1}{-1 + \frac{1}{-1+x}} = x$$

7.3 Palindromic Paths

Although nothing suggests it, the periodic path which we exhibited for $\sqrt{2}$ happens to be much more interesting than the one for the golden ratio. Notice that the second solution of the equation $x^2 - 2 = 0$ is $-\sqrt{2}$. In particular, the two solutions are opposites of each other. This fact alone will allow us to draw interesting conclusions about the periodic path and regular continued fraction for $\sqrt{2}$.

As before, we can write a continued fraction for the upper x using its ancestry of backward movements. But first, note that we can obtain in advance the continued fraction for $-\sqrt{2}$: in Section 5.2, we showed that multiplying a continued fraction by -1 is equivalent to making each of the coefficients negative. Thus, from a continued fraction for $\sqrt{2}$, we can directly write a continued fraction for $-\sqrt{2}$ by negating the coefficients.

$$x = 1 + \frac{1}{2 + \frac{1}{1+x}} \quad \rightarrow \quad -x = -\left(1 + \frac{1}{2 + \frac{1}{1+x}}\right) \quad \rightarrow \quad -x = -1 + \frac{1}{-2 + \frac{1}{-1-x}}$$

Now we have continued fractions which correspond with the ancestry of x in both directions. By construction, the two continued fractions have, up to negatives, the same coefficients. But those coefficients count the numbers of movements along the path in opposite directions. That is, when following the path forward or backward, the numbers of left and right movements are the same either way. So, the sequence of coefficients must be a palindrome!

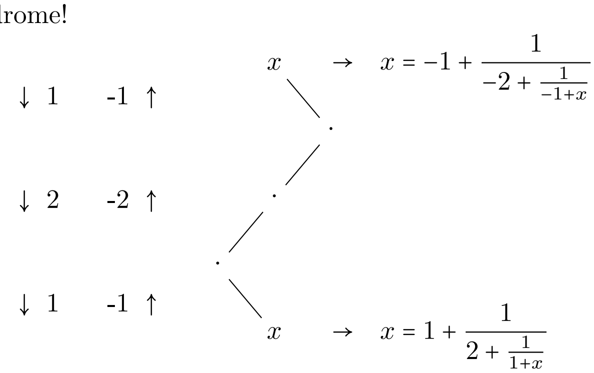

Of course, we can just recognize that 1, 2, 1 is a palindrome. But the argument is not specific to $\sqrt{2}$; we only used that the solutions to $x^2 - 2 = 0$ are symmetric. So, the same argument holds for all equations with symmetric roots, namely, those of the form $ax^2 - c = 0$ where a and c are positive integers such that $\sqrt{c/a}$ is not rational. Here are a few selected examples.

$$x = \sqrt{\frac{15}{7}} \qquad x = 1 + \cfrac{1}{2 + \cfrac{1}{6 + \cfrac{1}{2 + \cfrac{1}{1 + x}}}}$$

$$x = \sqrt{\frac{3}{85}} \qquad x = 0 + \cfrac{1}{5 + \cfrac{1}{3 + \cfrac{1}{5 + \cfrac{1}{0 + x}}}}$$

$$x = \sqrt{\frac{133}{24}} \qquad x = 2 + \cfrac{1}{2 + \cfrac{1}{1 + \cfrac{1}{4 + \cfrac{1}{1 + \cfrac{1}{2 + \cfrac{1}{2 + x}}}}}}$$

Look at the palindromes! Suppose in general that the sequence of coefficients is a_0, a_1, a_2, ..., a_2, a_1, a_0. By recursing, we find that the infinite continued fraction is given by $[a_0; \overline{a_1, a_2, \ldots, a_2, a_1, 2a_0}]$. This result is well-known, but its proof, even for the case of $\sqrt{n}$ where n is a positive integer, is long and demanding. One version of the proof is found in [16]. This Calkin-Wilf-inspired version is original and much more succinct.

7.4 Identity Paths

The two examples of periodic paths that we have seen so far may give a false impression that all regular periodic paths impose quadratic equations with two irrational solutions. This is not true; there are two families of periodic paths which do not have this property. They are the *all-left* and *all-right* periodic paths, containing only left movements or only right movements.

Applying the usual technique to solve for x, we find that the all-left periodic paths have equations of the form $mx^2 = 0$. Since m is positive, the only solution is $x = 0$. Likewise, the all-right periodic paths have equations of the form $m + x = x$. At first, this appears to have no solution, since $m \neq 0$. But in Section 6.3, we defined arithmetic with $1/0$ such that $m + 1/0 = 1/0$ for any m. Thus, the solution is $x = 1/0$. We see both left movement chains of $0/1$ and right movement chains of $1/0$ in the four-way tree in Section 6.1.

With these unimaginative paths covered, we can discuss a more surprising class of special periodic paths whose existence you might not expect. Consider the path below with forward and backward movements:

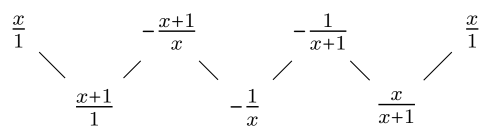

Can you tell what is unusual about this sequence? One end of the path is labelled $x/1$, and then the left and right rules are applied until the other end is reached. Remarkably, the final fraction is $x/1$ again! Now, when we set this equal to x, we get the equation $x = x$. This is not a quadratic; it is an identity which holds for all x. We call a periodic path with this property an *identity path*. These paths take each number to themselves. In this sense, identity paths capture certain translational symmetries of the four-way tree. In Section 8.3, we will give a more rigorous description. For now, we should at least appreciate that identity paths cannot be regular; they must feature a mix of forward and backward movements. That is because regular periodic paths have at most two solutions (all-left and all-right paths have one solution, while the rest have two irrational solutions), whereas identity paths are periodic paths for all values of x.

Identity paths come in all shapes and sizes. Here are a few more to consider:

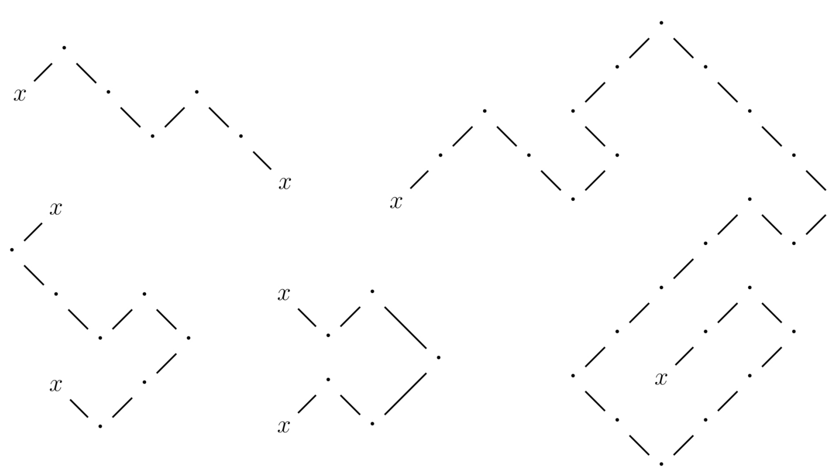

7.5 Factoring Paths and Oscillating Paths

Identity paths are the first indication that nonregular periodic paths are less well-behaved as regular ones. Our study of periodic paths began with the contrapositive statement that if a number were to appear more than once in the Calkin-Wilf tree (or double tree), then it is not rational. However, the same statement cannot be made for nonregular periodic paths, since *all* numbers appear more than once in the four-way tree! Because of this, we can construct nonregular periodic paths for any rational number, simply by locating that number twice in the four-way tree. The path between them may be an identity path, but this is not necessarily the case. For example, consider the nonregular periodic path below:

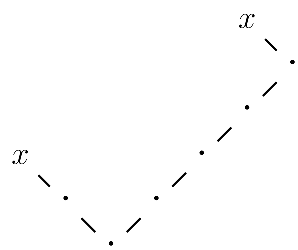

Following the left and right rules along the path and setting the other end equal to x, we get the equation $4x^2 + 12x + 9 = 0$. This equation factors as $(2x+3)^2 = 0$, so its only solution is $x = -2/3$. Indeed, this is a path from $-2/3$ to itself in the four-way tree.

How about another: this nonregular periodic path hides a different surprise.

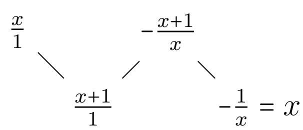

This time, the equation imposed on x is $x^2 + 1 = 0$, whose solutions are complex numbers! Notice, by the way, that this quadratic has symmetric roots, and its path is palindromic in accordance with Section 7.3.

By now, we have been conditioned to convert these paths into periodic continued fractions. But we should be cautious; as stated in Section 7.1, convergence is only guaranteed for regular continued fractions. Now that we are in four-way tree territory, the continued fractions corresponding with periodic paths may contain a mix of positive and negative coefficients, and therefore are not guaranteed to converge.

For example, the ancestry in the periodic path above yields the continued fraction $x = [1; -1, 1+x]$, which recurses to $x = [1; \overline{-1, 2}]$. To test for convergence, let's look at the convergents of this continued fraction.

$$
\begin{array}{cccccc}
1 & 1+\frac{1}{-1} & 1+\cfrac{1}{-1+\frac{1}{2}} & 1+\cfrac{1}{-1+\cfrac{1}{2+\frac{1}{-1}}} & 1+\cfrac{1}{-1+\cfrac{1}{2+\cfrac{1}{-1+\frac{1}{2}}}} & \cdots \\
1 & 0 & -1 & \frac{1}{0} & 1 & \cdots
\end{array}
$$

The convergents do not converge, but, rather, oscillate between the four values 1, 0, -1, and $1/0$. Accordingly, we refer to such a periodic path as an *oscillating path.* Without convergence, we cannot assign any value to $[1; \overline{-1, 2}]$, but we might hope that there exists some other context in which it makes sense to call this continued fraction $\sqrt{-1}$. This is left as an open research question. In any case, the nonregular periodic continued fractions which are produced by identity, factoring, and oscillating paths deserve further study.

8 Variation 4: Matrix Representation

8.1 Introduction

In Section 7, we became familiar with the constructions of periodic paths. Recall the workflow: we choose a sequence of left and right movements which comprise the path, and then, after applying the techniques of the section, we produce an irrational number and its corresponding infinite continued fraction. This method is simple and powerful, yet we have no control over the irrational numbers which are produced. That is, if one is interested in a periodic path for, say, $\sqrt[3]{2}$, we have not specified an approach for constructing such a path. In fact, the continued fraction for $\sqrt[3]{2}$ may not be periodic, so this task is not always possible in general. Our new goal should be to understand exactly which irrational numbers have regular periodic paths. As we will see, the answer to this question lies in a new representation of the Calkin-Wilf tree using matrices.

To begin the translation, we must pass from fractions to vectors. Naturally, we represent the fraction p/q with the column vector $\left(\begin{smallmatrix} p \\ q \end{smallmatrix}\right)$. Of course, there are many ways to represent fractions, including ordered pairs, but we choose vectors because they can be operated on by matrices. In particular, we would like to define a left rule matrix L such that $L\left(\begin{smallmatrix} p \\ q \end{smallmatrix}\right) = \left(\begin{smallmatrix} p \\ p+q \end{smallmatrix}\right)$. Simply, $L = \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$. Likewise, the matrix $R = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ performs the right rule; $R\left(\begin{smallmatrix} p \\ q \end{smallmatrix}\right) = \left(\begin{smallmatrix} p+q \\ q \end{smallmatrix}\right)$. These same observations are made in [15]. We will call these two special matrices the *left matrix* and the *right matrix*, respectively. Notice that the left and right matrices satisfy $L^m = \left(\begin{smallmatrix} 1 & 0 \\ m & 1 \end{smallmatrix}\right)$ and $R^m = \left(\begin{smallmatrix} 1 & m \\ 0 & 1 \end{smallmatrix}\right)$. Thus, we can make multiple movements through the vectorized Calkin-Wilf tree using L^m and R^m as analogues of the m-left and m-right rules.

So far, the translation is going well. But with the goal to understand periodic paths, we must find a matrix representation of those as well. For terminology, we will call a path which starts at $x/1$ a *relative path*, since the fractions along its length are all expressed relative to value of x. Let's pay closer attention to the fractions which appear along relative paths. Since the left and right rules add numerators to denominators and vice versa, it can only be that the numerators and denominators contain multiples of x and 1. So, in general, fractions along a relative path have the form $(ax+b)/(cx+d)$ for integers a, b, c, and d. We choose to represent such a fraction by the matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, since $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)\left(\begin{smallmatrix} x \\ 1 \end{smallmatrix}\right) = \left(\begin{smallmatrix} ax+b \\ cx+d \end{smallmatrix}\right)$.

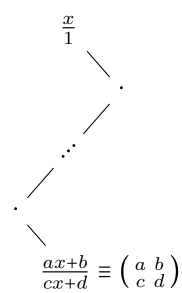

Importantly, the fraction $x/1$ corresponds with the identity matrix $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ in this representation. This means that all fractions along relative paths correspond with matrices which are purely products of L and R. For example, the product below computes the fraction at the end of a relative path containing 2 right, 3 left, and 1 right movement. Setting $x = 0/1$, the relative path becomes a path in the Calkin-Wilf tree; this one is the path from 0/1 to 9/7, which we saw in Section 4.1.

$$R^1 L^3 R^2 = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & 0 \\ 3 & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & 2 \\ 0 & 1 \end{smallmatrix}\right) = \left(\begin{smallmatrix} 4 & 9 \\ 3 & 7 \end{smallmatrix}\right) \equiv \frac{4x+9}{3x+7} \xrightarrow{x=0} \frac{9}{7}$$

This suggests something more interesting: we can use relative paths to generalize the Calkin-Wilf tree. If, at any time, we would like to return to the Calkin-Wilf tree, we just set $x = 0/1$. To be concrete, here is a diagram of what we call the *relative Calkin-Wilf tree* (also called the 'Calkin-Wilf tree of x' [14]), along with its corresponding matrix representation, which we call the *matrix tree*. Paths in matrix tree are traversed by left-multiplying nodes by the left and right matrices. The Calkin-Wilf tree is obtained from the matrix tree by considering only the second column of each matrix.

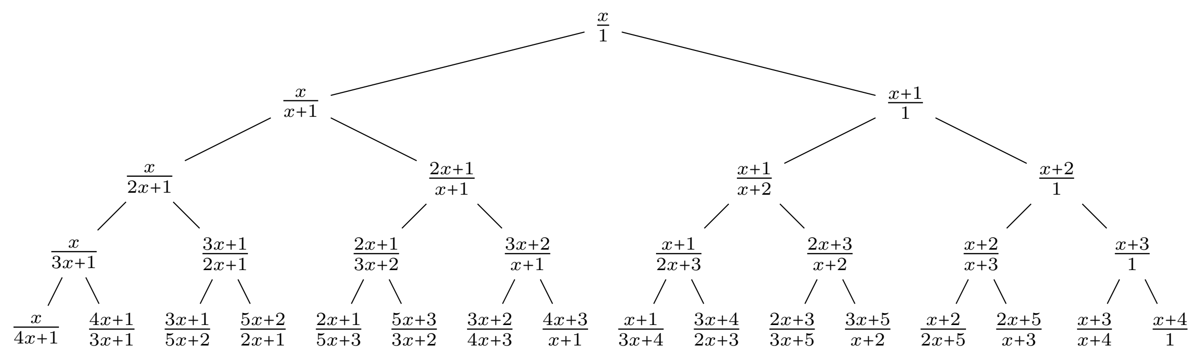

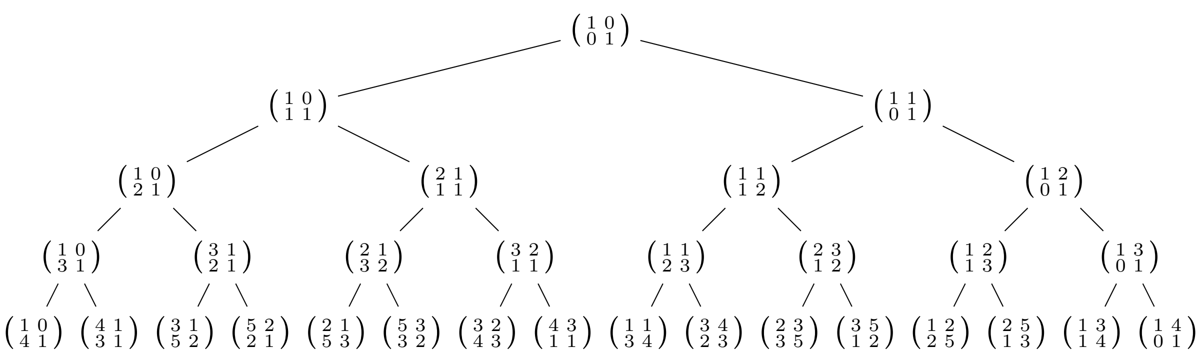

Before pinning down exactly which matrices appear in this tree, we will make a brief remark about the Euclidean algorithm.

8.2 Euclidean Algorithm in the Matrix Tree

Recall from Section 4.1 that the Euclidean algorithm and the Calkin-Wilf tree are closely related: backward paths through the tree perform the slow Euclidean algorithm, and the fractions at the turning points along that path contain parts of the intermediate computations of the fast Euclidean algorithm. We can use this correspondence to derive a similar algorithm for computing the ancestry of a node in the matrix tree.

Our goal is to retrace the ancestry of $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. We begin from the observation in Section 8.1 that the second columns of matrices in the matrix tree are the vector representations of fractions in the Calkin-Wilf tree. Thus, ignoring the first column, we can perform the Euclidean algorithm on the second column of $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ to obtain a continued fraction for b/d. Remember to adjust the continued fraction if necessary so that it has an odd number of coefficients. In the example below with the matrix $\left(\begin{smallmatrix} 3 & 1 \\ 5 & 2 \end{smallmatrix}\right)$, the Euclidean algorithm admits for $1/2$ the continued fraction $[0; 2] = [0; 1, 1]$. This corresponds with a path in the matrix tree to $\left(\begin{smallmatrix} 3 & 1 \\ 5 & 2 \end{smallmatrix}\right)$ from some matrix whose second column is $\left(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix}\right)$. To obtain the first column of that matrix, we just perform Euclid's algorithm on the first column of $\left(\begin{smallmatrix} 3 & 1 \\ 5 & 2 \end{smallmatrix}\right)$ in tandem with the second column. As a result, this matrix will have the form $\left(\begin{smallmatrix} 1 & 0 \\ m & 1 \end{smallmatrix}\right)$ (we will justify this next). In the example below, the matrix is $\left(\begin{smallmatrix} 1 & 0 \\ 2 & 1 \end{smallmatrix}\right)$. But this matrix is exactly L^2, or L^m in general. This completes the decomposition of $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ into a product over L and R.

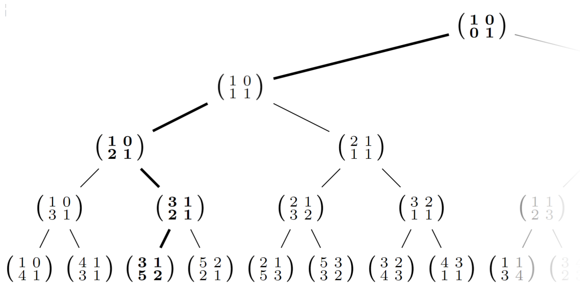

$$\frac{1}{2} = 0 + \cfrac{1}{1 + \frac{1}{1}}$$

$$\begin{pmatrix} 3 & 1 \\ 5 & 2 \end{pmatrix} = R^0 L^1 R^1 \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} = L^1 R^1 L^2$$

The missing detail is resolved from the observation that L and R both have determinant 1. So, by the product rule for determinants, every node in the matrix tree has determinant 1. Therefore, if the second column of $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is $0/1$, then the determinant condition forces $a = 1$, as we claimed. In summary, this application of the Euclidean algorithm shows that each $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ in the matrix version of the relative Calkin-Wilf tree has a unique decomposition as a product over L and R. This result also appears in [15].

8.3 Classifications

Finally, we are ready to learn which irrational numbers correspond with regular periodic paths. We already know that relative paths contain fractions of the form $(ax + b)/(cx + d)$, which we choose to represent with the matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$. It is left to determine the possible values of a, b, c, and d. When the relative path is regular (consisting of only forward movements), we have from Section 8.2 that $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ can be expressed as a product over L and R. So, we only need to know what is generated by L and R under multiplication. In fact, it is exactly $SL_2(\mathbb{N})$, which the set of 2 by 2 matrices with coefficients in $\mathbb{N}$ and determinant 1.

Now, we're only one step away. Since numbers with regular periodic paths are obtained by forcing some $(ax + b)/(cx + d)$ at the end of a relative path equal to x, we conclude that there is a regular periodic path for x exactly when $x = (ax + b)/(cx + d)$ for some $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ in $SL_2(\mathbb{N})$. That's the characterization! For example, in Section 7.2, we found that the golden ratio satisfies $x = (2x + 1)/(x + 1) \equiv \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$, and $\sqrt{2}$ satisfies $x = (3x + 4)/(2x + 3) \equiv \begin{pmatrix} 3 & 4 \\ 2 & 3 \end{pmatrix}$; both of these matrices have determinant 1.

As a corollary, since regular periodic paths correspond with regular periodic continued fractions (which are guaranteed to converge), we conclude that the numbers x satisfying the condition above are exactly those with periodic continued fractions.

What about nonregular periodic paths? The characterization is boring, since we noted in Section 7.5 that there are such paths for all rational numbers. We even found a nonregular periodic path for $\sqrt{-1}$. While it is harmless to establish these paths, we must be cautious when talking about the periodic continued fractions that they produce, since these continued fractions are not guaranteed to converge.

On the subject of nonregularity, we should mention that it is also possible to construct a relative four-way tree. There is a corresponding four-way matrix tree, and the matrices which appear in it are generated by L, R, and their inverses. We find that $L^{-1} = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}$ and $R^{-1} = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}$. Consequently, $L^{-m} = \begin{pmatrix} 1 & 0 \\ -m & 1 \end{pmatrix}$ and $R^{-m} = \begin{pmatrix} 1 & -m \\ 0 & 1 \end{pmatrix}$, which is reminiscent of the m-backward-left and right rules having interpretations as $(-m)$-left and right rules (see Section 6.2). Under multiplication, L, R, and their inverses generate $SL_2(\mathbb{Z})$, the special linear group. It is possible to establish a Euclidean-like algorithm which traces paths in the four-way matrix tree; the catch is that the entries of a matrix in $SL_2(\mathbb{Z})$ may have different signs. Since ancestry is not unique, the algorithm can be used to generate identity paths like the sophisticated one shown in Section 7.4. All identity paths correspond with *relations* (sequences whose products are $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$) over L, R, and their inverses.

9 Further Studies

The topics of listing the rationals, continued fractions, and $SL_2(\mathbb{Z})$ share plenty of related content. Below, we briefly describe a few areas of further study which are tangential to these topics.

9.1 The Stern-Brocot Tree

Related to the Calkin-Wilf tree is the Stern-Brocot tree, another tree which contains each positive rational number exactly once [7]. A child in the Stern-Brocot tree is obtained by taking the *mediant* of its most recent left and right ancestors: if those two ancestors are a/b and c/d, then the child is $(a+b)/(c+d)$. A depth-first search of the tree produces a list of rational numbers ordered by magnitude. In both the Calkin-Wilf and Stern-Brocot trees, the root 1/1 is the right child of 0/1 and the left child of 1/0. The two trees are shown below for comparison, with the Calkin-Wilf tree on the left and the Stern-Brocot tree on the right.

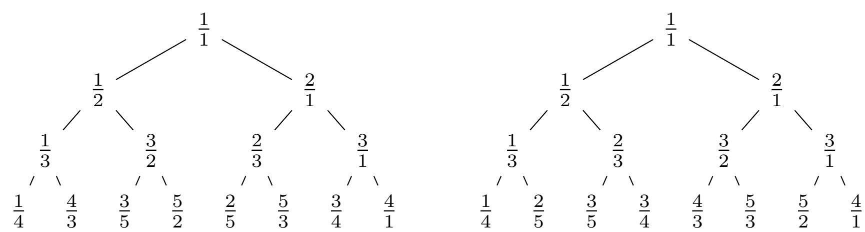

The two trees may appear identical at a glance. In fact, the set of fractions in each level is the same between both trees. In other words, one tree can be transformed into the other by shuffling the fractions in each level. The particular permutation is constructed as follows: in level n, label the entries 0 to $(n-1)$, and convert the labels to binary. Then, reverse the order of the bits. Converting back to decimal, this establishes a permutation on the entries which rearranges level n of one tree into the same level of the other [2]. Note that bit reversal is an order-2 action, so each permutation is a disjoint product of 2-cycles. Thus, these level permutations turn one tree into the other, and vice versa. For example, here is the permutation which transforms the third level of one tree into the third level of the other tree. In cycle notation, this is (14)(36).

$\frac{1}{4}$	$\frac{4}{3}$	$\frac{3}{5}$	$\frac{5}{2}$	$\frac{2}{5}$	$\frac{5}{3}$	$\frac{3}{4}$	$\frac{4}{1}$
0	1	2	3	4	5	6	7
000	001	010	011	100	101	110	111
000	100	010	110	001	101	011	111
0	4	2	6	1	5	3	7
$\frac{1}{4}$	$\frac{2}{5}$	$\frac{3}{5}$	$\frac{3}{4}$	$\frac{4}{3}$	$\frac{5}{3}$	$\frac{5}{2}$	$\frac{4}{1}$

Like the Calkin-Wilf tree, paths in the Stern-Brocot tree are related to continued fractions, and the Stern-Brocot tree also has a matrix representation for its left and right rules; details can be found in [7].

The *Stern-Brocot sequences of order* n are obtained by performing a depth-first search on the first n levels of the tree. These sequences have commonalities with the *Farey sequences*, which are also obtained by taking mediants of fractions. Farey sequences are related to *Ford circles*, which, in turn, are related to linear fractional transformations (see Section 9.3) and $SL_2(\mathbb{Z})$. All of these topics, along with continued fractions, are highly interconnected. Many of the relevant details can be found by an online search.

9.2 The Hyperbinary Sequence

Take a closer look at the Calkin-Wilf sequence: $\{1/1, 1/2, 2/1, 1/3, 3/2, 2/3, 3/1, 1/4, \dots\}$. Does anything stand out? We haven't yet appreciated how consecutive terms are related. The pattern to notice is that the denominator of one fraction is the numerator of the next! In other words, the sequence of numerators is exactly the same as the sequence of denominators, shifted by one place. This sequence, $\{1, 1, 2, 1, 3, 2, 3, 1, 4, \dots\}$, is called the *hyperbinary sequence*, denoted $b(n)$ for $n \geq 0$. Each term, n, counts the number of ways to write n in binary where the place values are used at most twice (in ordinary binary, the place values are used at most once). For example, $b(4)$ counts the hyperbinary representations of four: 12, 20, and 100, which are $(2+1+1)$, $(2+2)$, and (4), respectively.

Terms in the Calkin-Wilf sequence are ratios of consecutive terms in the hyperbinary sequence. In particular, $\ell(n) = b(n-1)/b(n)$. This is how the Calkin-Wilf sequence was introduced by Calkin and Wilf in their expository article [5].

In Section 2.2, we noted two fundamental properties of the Calkin-Wilf sequence. From both, we can draw conclusions about hyperbinary sequence. First, since each fraction in $\ell(n)$ is reduced, it follows that consecutive terms in $b(n)$ are coprime. Second, since every possible reduced rational appears exactly once in $\ell(n)$, we have that every possible ordered pair of integers appears exactly once as a pair of consecutive terms in $b(n)$. For example, there is exactly one instance of a 9 followed by a 7 in $b(n)$.

Additionally, the structure of the hyperbinary sequence is described by left and right rules. If a term in the Calkin-Wilf tree has numerator $b(n)$, then its left and right children have numerators $b(2n+1)$ and $b(2n+2)$. But the left and right children of $b(n)/b(n+1)$ have numerators $b(n)$ and $(b(n)+b(n+1))$ using the left and right rules. Comparing expressions, we obtain the following recursive rules for $b(n)$:

$$b(2n+1) = b(n)$$
$$b(2n+2) = b(n) + b(n+1)$$

Together with $b(0) = 1$, these rules generate the entire sequence $b(n)$.

To learn more about the hyperbinary sequence, one may begin with the references under A002487 on the Online Encyclopedia of Integer Sequences.

9.3 Linear Fractional Transformations

A *linear fractional transformation* (LFT) is a function the form

$$f(z) = \frac{az+b}{cz+d} \qquad ad - bc \neq 0 \tag{1}$$

where the coefficients a, b, c, and d are in $\mathbb{C}$. The name 'linear fractional' references the transformation's algebraic structure: a rational function whose numerator and denominator are linear functions. In some texts, LFTs are also called 'bilinear transformations' and 'Möbius transformations' [18]. Note that if $c = 0$, then f is a linear transformation. Otherwise, if $c \neq 0$, then the denominator $cz + d$ approaches ∞ as z approaches $-d/c$. Thus, it is common to consider LFTs over the extended complex numbers $\widehat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$, where we set $f(\infty) = b/d$ if $c \neq 0$ and $f(\infty) = \infty$ otherwise.

We will refer to the condition $ad - bc \neq 0$ as the *non-degenerate condition*, since when $ad - bc = 0$, it follows that f is a constant function. In contrast, when the non-degenerate condition is met, f is bijective over $\widehat{\mathbb{C}}$ and, therefore, is invertible. Solving (1) for z, we find that f^{-1} is also an LFT.

$$f^{-1}(z) = \frac{dz - b}{-cz + a}$$

The composition of two LFTs is also an LFT; below is the composition of $f(z) = \frac{az+b}{cz+d}$ and $g(z) = \frac{jz+k}{lz+m}$.

$$f(g(z)) = \frac{a\left(\frac{jz+k}{lz+m}\right) + b}{c\left(\frac{jz+k}{lz+m}\right) + d} = \frac{ajz + ak + blz + bm}{cjz + ck + dlz + dm} = \frac{(aj + bl)z + (ak + bm)}{(cj + dl)z + (ck + dm)}$$

Both the formulations of composition and inversion are suggestive of a beautiful idea: if we represent each LFT in the form of (1) with the 2 by 2 matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, then LFT composition corresponds with matrix multiplication, and LFT inversion corresponds with matrix inversion (this is always possible due to the non-degenerate condition). Moreover, the identity function $f(z) = z$ is an LFT, and it is represented by the identity matrix $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$.

In general, the coefficients of an LFT can be scaled together by a nonzero number without changing the LFT. That is, for any $s \neq 0$, the LFT with coefficients a, b, c, and d is equivalent to the LFT with coefficients sa, sb, sc, and sd. In matrix form, $s \cdot \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ represents the same LFT as $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. This explains why the inverse of an LFT is represented by the matrix $\left(\begin{smallmatrix} d & -b \\ -c & a \end{smallmatrix}\right)$: we can scale $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)^{-1}$ by $s = \det\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, which is nonzero by the non-degenerate condition.

Recall that the *fixed points* of a function f are the values of z such that $f(z) = z$. If f is an LFT, its fixed points satisfy $(az + b)/(cz + d) = z$. When $c \neq 0$, we can write the fixed points of f in terms of a, b, c, and d using the quadratic formula. Otherwise, if $c = 0$, then $f(z)$ fixes ∞. When f is the identity function, every z in the domain is fixed.

Do these formulations look familiar? In fact, LFTs offer another interpretation of relative and periodic paths. Recall from Section 8.3 that the fractions along a relative path have the form $(ax + b)/(cx + d)$, where $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ is in $SL_2(\mathbb{Z})$. In the context of LFTs, we may think of a relative path as a function, mapping its first node $x/1$ to its last node $(ax + b)/(cx + d)$. This function is exactly the LFT $f(x) = (ax + b)/(cx + d)$.

We turn a relative path into a periodic path for x by setting the other end of the path equal to x. The values of x are found by solving $(ax + b)/(cx + d) = x$. Thus, x has a certain periodic path exactly when it is a fixed point of the LFT corresponding with that periodic path.

Of course, we cannot ignore that both relative paths and LFTs have matrix representations. By construction, relative paths and their corresponding LFTs have the same matrix representations. So, with matrices as a link, it is possible to translate properties of LFTs to relative paths. This is convenient since LFTs and their properties are well-documented. For an introduction, one may begin with [4].

9.4 Miscellaneous Continued Fractions

Continued fractions are of current interest: a search on MathSciNet for titles with the keyword 'continued fraction' returns hundreds of articles from recent years, 19 of which appeared in 2023.

Recall that paths in the Calkin-Wilf tree correspond with regular continued fractions, and paths in the four-way tree correspond with continued fractions whose coefficients are integers. These are not the only types of continued fractions. In fact, many authors have considered continued fractions with coefficients from other sets, including the Gaussian integers (see *Hurwitz continued fractions*) [17], the Heisenberg group [11], Iwasawa groups [12], p-adic fields [9], and square matrices over $\mathbb{R}$ and $\mathbb{C}$ [13]. Certainly, there is plenty to explore if one is interested in continued fractions and their applications to various mathematical fields.

References

[1] H. Appelgate and H. Onishi. Continued Fractions and the Conjugacy Problem in $\mathrm{SL}_2(\mathrm{Z})$. *Communications in Algebra*, 9(11):1121–1130, 1981.

[2] B. Bates, M. Bunder, and K. Tognetti. Linking the Calkin–Wilf and Stern–Brocot Trees. *European Journal of Combinatorics*, 31(7):1637–1661, 2010.

[3] J. Berstel and A. de Luca. Sturmian Words, Lyndon Words and Trees. *Theoretical Computer Science*, 178(1-2):171–203, 1997.

[4] J.W. Brown and R.V. Churchill. *Complex Variables and Applications*. McGraw-Hill, 2009.

[5] N. Calkin and H.S. Wilf. Recounting the Rationals. *The American Mathematical Monthly*, 107(4):360–363, 2000.

[6] B. Gobler. Listing the Rationals using Continued Fractions. *The Pi Mu Epsilon Journal*, 15(6):347–354, 2022.

[7] R. Graham, D. Knuth, and O. Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. A Foundation for Computer Science. Addison-Wesley, 1989.

[8] J.E. Graver. Listing the Positive Rationals. *Mathematics Magazine*, 94(1):24–33, 2021.

[9] J. Hirsh and L.C. Washington. P-adic Continued Fractions. *The Ramanujan Journal*, 25:389–403, 2011.

[10] O. Karpenkov. *Geometry of Continued Fractions*. Algorithms and Computation in Mathematics. Springer Berlin Heidelberg, 2013.

[11] A. Lukyanenko and J. Vandehey. Continued Fractions on the Heisenberg Group. *Acta Arithmetica*, 1(167):19–42, 2015.

[12] A. Lukyanenko and J. Vandehey. Ergodicity of Iwasawa Continued Fractions via Markable Hyperbolic Geodesics. *Ergodic Theory and Dynamical Systems*, 43(5):1666–1711, 2023.

[13] S. Mennou, A. Chillali, and A. Kacha. Matrix Continued Fractions and Expansions of the Error Function. *arXiv preprint arXiv:2211.16453*, 2022.

[14] L. Ponton. The Calkin-Wilf Tree of a Quadratic Surd. *The American Mathematical Monthly*, 126(9):771–785, 2019.

[15] G.N. Raney. On Continued Fractions and Finite Automata. *Mathematische Annalen*, 206:265–283, 1973.

[16] K.H. Rosen. *Elementary Number Theory*. Pearson Education London, 2011.

[17] D. Simmons. The Hurwitz Continued Fraction Expansion as Applied to Real Numbers. *L'Enseignement Mathématique*, 62(3):475–485, 2017.

[18] D.C. Ullrich. *Complex Made Simple*. Prilog, 2008.