The Dynamics of Continued Fractions

Evan O'Dorney

May 31, 2011

1 The Story

I was first introduced to the Intel Science Talent Search in ninth grade. I knew I would have no trouble entering this contest, as I had just completed a research project in linear algebra that year. I also knew I would not have to think about the contest for a long while, as it is only open to twelfth graders.

During the summer, I received an invitation from Brian Conrad, a professor at Stanford University, to discuss a recently formulated research problem in number theory. Conrad had heard of me from fellow professor Ravi Vakil, who in turn had witnessed my passion for mathematics at the Berkeley Math Circle. In the past, Conrad had invited various students to explore the problem; the latest results, proved by Shravani Mikkilineni [1], had won fourth prize at the Intel Science Talent Search. However, Mikkilineni had made clear that her work was not the end of the story by formulating a *conjecture*—a precise statement of a result not yet proved. It was this conjecture that Conrad presented to me at Stanford.

I love working on conjectures. Just as in the various Mathematical Olympiads in which I have participated, the conditions are already set; the challenge consists in cleverly using the hypotheses of the problem to produce the conjectured conclusion. Number theory, the study of properties of the ordinary counting numbers $1, 2, 3, \ldots$, is particularly rich in this type of problems, which range from puzzles for the general audience to the challenges on the International Mathematical Olympiad to famous conjectures, such as Fermat's Last Theorem and the Twin Prime Conjecture, which commonly remain unsolved for hundreds of years. My instincts told me that Conrad's problem would not be one of these enduring conjectures, and eagerly I set to work.

Unfortunately, school was about to start, and I made little progress other than simplifying some notations that I found more complex than necessary. The following summer I began working again, using the computer system *Mathematica* which I had won through my performance on the Mathematical Olympiad. After generating and analyzing copious data for two weeks, I could finally write down the key formula that enabled me to prove Mikkilineni's conjecture.

My mathematical investigations were still not finished. For weeks thereafter, I wrote down my proofs, carefully searching for ways to simplify them. I found that the methods that I had invented would work in far more general circumstances than Mikkilineni had suggested—and I could describe exactly what these circumstances were. Months later, while preparing a Math Circle talk on this project, which I thought was complete, I discovered to my surprise that the main condition could be simplified even further (giving the form (1) below). By this time, on seeing my results, my professors unanimously opined that this number theory project was much stronger than the linear algebra project that I had originally planned to submit to Intel.

2 The Mathematics

Every high school student has seen $\pi = 3.14159\ldots$, the ratio of a circle's circumference to its diameter, and its two widely used approximate values: 3.14

and 22/7. The value 3.14 clearly arises from dropping all but the first two digits after the decimal point in the infinite decimal form $3.14159\ldots$, but 22/7 is somewhat a mystery: where do 22 and 7 come from? They are clearly not random; the quotient $3.1428\ldots$ is within two thousandths of π, much closer than we would expect from a fraction with denominator only 7.

In general, we have the following problem:

Question 1. If x is an irrational number (one that cannot be written exactly as a fraction), how can we find the "extraordinary" approximations a/b where a and b are natural numbers as small as possible while a/b is as close to x as possible?

For the past 200 years, this question has had a standard answer—the method of *continued fractions*. To explain how it works, suppose we are trying to approximate π. We notice that π is "3 plus something," and we try to write that "something" as a fraction with 1 as its numerator:

$$\pi = 3.14159\ldots = 3 + 0.14159\ldots = 3 + \frac{1}{7.06251\ldots}$$

Here, we can notice that the denominator is very close to 7. Therefore $\pi \approx 3 + \frac{1}{7}$ which, when simplified, yields the familiar approximation 22/7.

Or, we can treat the denominator as "7 plus something" and repeat the process:

$$3 + \frac{1}{7.06251\ldots} = 3 + \frac{1}{7 + 0.06251\ldots} = 3 + \cfrac{1}{7 + \cfrac{1}{15.99659\ldots}}$$

If we approximate the last denominator by 16, we get

$$\pi \approx 3 + \cfrac{1}{7 + \cfrac{1}{16}} \quad \text{which simplifies to} \quad \frac{355}{113}.$$

This is a remarkable approximation to π—accurate to 6 decimal places—first discovered by the Chinese mathematician Zu Chongzhi in the fifth century.

But it is still possible to go on, and in the end we get an infinite *continued fraction* for π:

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

To save space, mathematicians prefer to write the terms horizontally:

$$\pi = [3, 7, 15, 1, 292, 1, 1, 1, 2, 1, 3, 1, 14, 2, 1, 1, 2, \ldots].$$

Stopping this continued fraction at any point gives a fraction that approximates π. The approximations converge on the true value of π quite fast:

$$\begin{aligned}
[3] &= 3 = 3 \\
[3,7] &= \frac{22}{7} = 3.142857142857\ldots \\
[3,7,15] &= \frac{333}{106} = 3.141509433962\ldots \\
[3,7,15,1] &= \frac{355}{113} = 3.141592920353\ldots \\
[3,7,15,1,292] &= \frac{103993}{33102} = 3.141592653011\ldots \\
[3,7,15,1,292,1] &= \frac{104348}{33215} = 3.141592653921\ldots \\
\pi &= 3.14159265358979\ldots
\end{aligned}$$

Moreover, the approximations to an irrational number x produced by the continued fraction method are the "best" answers to Question 1 in the following sense [2]:

Theorem. *Let a/b be an approximation to an irrational number x found by the continued fraction method. Let c/d be any fraction such that c and d are natural numbers and $d < b$. Then a/b is closer to x than c/d is.* (In other words, a/b sets a record of accuracy among fractions with a given denominator size.)

The continued fraction method can be applied to any irrational number and, as can be seen from the π example, usually yields very random terms. Beautiful patterns appear, however, if we consider not π but a *quadratic* irrational, the square root of a positive non-square integer.

$$\sqrt{3} = [1, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, \ldots]$$

$$\sqrt{45} = [6, 1, 2, 2, 2, 1, 12, 1, 2, 2, 2, 1, 12, \ldots]$$

$$\sqrt{46} = [6, 1, 3, 1, 1, 2, 6, 2, 1, 1, 3, 1, 12, 1, 3, 1, 1, \ldots]$$

The following patterns may be noticed:

1. All of these continued fractions are *periodic*, having a section of numbers that repeats over and over again.

2. The only number not part of the repeating pattern is the first number (green).

3. The last number of the repeating pattern (red) is also the largest number, and it is twice as big as the first number (green).

4. The terms between these colored numbers (black) have a *palindromic* pattern, one that reads the same backwards and forwards.

The method of continued fractions described here may be aptly termed the "Classic Method" for answering Question 1; the most advanced of the properties described—the repeating pattern in the terms for a square root—was proved by Lagrange in 1770. By contrast, the "Dynamic Method," which I also analyzed in my Intel project, has only been studied as a solution to Question 1 for the past five years, though it too has roots dating back to antiquity.

Suppose we want to approximate $\sqrt{45}$. (The Dynamic Method does not apply to non-quadratic irrationals like π.) Consider this formula:

$$f(x) = \frac{6x + 45}{x + 6}$$

Here, 45 is clearly coming from the number whose square root we want. The two 6's are both coming from the *integer part* of the square root; here $\sqrt{45} = 6.70820\ldots$, and the integer part (the part before the decimal point) is 6.

The f in this method is a *function*: we can put any number in for x and get a number out. We start with 6, the previously determined integer part of the square root, and compute

$$f(6) = \frac{6 \cdot 6 + 45}{6 + 6} = \frac{27}{4}.$$

We then treat f as a dynamical system (this is the origin of the term Dynamic Method) by feeding the output back in as the input:

$$f\left(\frac{27}{4}\right) = \frac{114}{17}$$
$$f\left(\frac{114}{17}\right) = \text{etc.}$$

In this way, we get an infinite sequence of fractions:

$$6 \quad \frac{27}{4} \quad \frac{114}{17} \quad \frac{161}{24} \quad \cdots$$

It is not obvious from the definition of the function f whether these fractions converge to any limit. But given that they do, the limit must be a number that, plugged into f, gives the same number again, i.e. a solution of

$$x = \frac{6x + 45}{x + 6}.$$

Students with training in algebra are invited to solve this equation and find the two roots:

$$x = \sqrt{45} \quad \text{and} \quad x = -\sqrt{45}.$$

Because all the fractions are positive, only the positive root is relevant, and we may conclude (though not rigorously, since we have not proved that the limit exists) that this method yields a sequence of rational approximations to $\sqrt{45}$.

This process, described for $\sqrt{45}$, will work in general: given any non-square integer k, with d the integer part of $\sqrt{k}$, repeated application of the function

$$f(x) = \frac{dx + k}{x + d},$$

with starting point $x = d$, yields a sequence of fractions that invariably converge to $\sqrt{k}$. Graphically, this method is shown in Figure 1: the spiral line shuttling between the curves $y = x$ and $y = f(x)$ converges on their intersection, the red point $(\sqrt{k}, \sqrt{k})$.

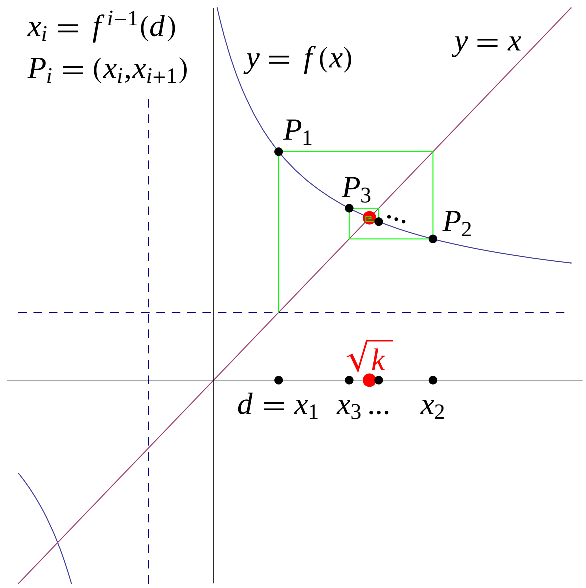

Figure 1: $f^n(d) \to \sqrt{k}$ as $n \to \infty$

A distinctive feature of the Dynamic Method is that it is a fast computation. This is probably evident from the definition, which involves merely applying the same function f again and again, but the speed can be increased further using a technique known as *repeated squaring*. To explain this, suppose we want to compute 2^{100}. Literally, this means multiplying 2 by itself 100 times, entailing 99 multiplications. But there are faster ways: if we know 2^{50}, we can square it to get $2^{50 \cdot 2} = 2^{100}$. Similarly, if we know 2^{25}, we can square it to get 2^{50}, and if we know 2^{12}, we can square (thus getting 2^{24}) and then multiply by 2 to get 2^{25}. Continuing in this way, we obtain a procedure for building 2^{100} up from 2^1:

$$2^{100} = \left(\left(\left(\left(2^2 \cdot 2\right)^2\right)^2 \cdot 2\right)^2\right)^2 .$$

Thus, 7 multiplications (5 squarings and 2 doublings) suffice to compute 2^{100}. The iterations of the Dynamic Method are amenable to a similar treatment, and using 7 operations (each representing combining a function either with itself or with f) it is possible to extract the 100th approximation to a given square root. This contrasts with the Classic Method, in which computing the 100th approximation entails finding 100 terms of the continued fraction, one by one.

If the Dynamic Method is very fast, the question arises of whether it is also accurate. Specifically:

Question 2. Does the Dynamic Method yield infinitely many of the extraordinary approximations to a square root, as determined by the Classic Method?

That a sequence of approximations *converges* is not enough for them to be *extraordinary* as desired in Question 1. For instance, the fractions

$$3, \frac{31}{10} = 3.1, \frac{314}{100} = 3.14, \frac{3141}{1000} = 3.141, \ldots,$$

derived from the decimal expansion of π, clearly *converge* to π, but their accuracy is quite ordinary for their denominator size, and they have nothing of the succinctness and exceptional accuracy of 22/7 and 355/113.

With the Dynamic Method, the quality of the approximations varies widely depending on the particular square root, as shown in the following table:

k	Approximations to $\sqrt{k}$ (Classic Method and Dynamic Method)
42	6 $\frac{13}{2}$ $\frac{162}{25}$ $\frac{337}{52}$ $\frac{4206}{649}$ $\frac{8749}{1350}$ $\frac{109194}{16849}$ $\frac{227137}{35048}$ $\frac{2834838}{437425}$ $\cdots$
	6 $\frac{13}{2}$ $\frac{162}{25}$ $\frac{337}{52}$ $\frac{4206}{649}$ $\frac{8749}{1350}$ $\frac{109194}{16849}$ $\frac{227137}{35048}$ $\frac{2834838}{437425}$ $\cdots$
43	6 7 $\frac{13}{2}$ $\frac{46}{7}$ $\frac{59}{9}$ $\frac{341}{52}$ $\frac{400}{61}$ $\frac{1541}{235}$ $\frac{1941}{296}$ $\frac{3482}{531}$ $\frac{43725}{6668}$ $\frac{47207}{7199}$ $\cdots$
	6 $\frac{79}{12}$ $\frac{990}{151}$ $\frac{12433}{1896}$ $\frac{156126}{23809}$ $\frac{1960543}{298980}$ $\frac{24619398}{3754423}$ $\frac{309156577}{47145936}$ $\cdots$
45	6 7 $\frac{20}{3}$ $\frac{47}{7}$ $\frac{114}{17}$ $\frac{161}{24}$ $\frac{2046}{305}$ $\frac{2207}{329}$ $\frac{6460}{963}$ $\frac{15127}{2255}$ $\frac{36714}{5473}$ $\frac{51841}{7728}$ $\frac{658806}{98209}$ $\cdots$
	6 $\frac{27}{4}$ $\frac{114}{17}$ $\frac{161}{24}$ $\frac{2046}{305}$ $\frac{8667}{1292}$ $\frac{36714}{5473}$ $\frac{51841}{7728}$ $\frac{658806}{98209}$ $\cdots$

The simplest type of behavior is shown in the first row, for $k = 42$: The Dynamic Method simply duplicates the results of the Classic Method. The conditions for this behavior to occur were discovered and proved by Rosen, Shankar, and Thomas in 2006 [3]: it happens if and only if

$$\frac{2d}{k - d^2}$$

is an integer.

In contrast to this is the behavior for $k = 43$. Here none of the fractions in the two rows match except for the trivial first approximation, 6. Moreover, the numerators and denominators for the Dynamic Method grow much faster than in any other row of the table, highlighting the fact that they are accurate without being extraordinary.

My research focused on explaining the type of behavior shown in the row $k = 45$. Some but not all of the Classic and Dynamic approximations coincide, and moreover they form a peculiar pattern: the matching terms (yellow) occur in groups of three, separated by three dissimilar terms in the Classic sequence and one term in the Dynamic sequence. Discovering which square roots behave in this way was not easy, but in the end I found a formula similar to Rosen, Shankar, and Thomas's: if

$$\frac{4d^2}{k-d^2} \tag{1}$$

is an integer, then infinitely many of the Classic and Dynamic approximations match, and there is always a repeating pattern similar to that for $k = 45$. If (1) is not an integer, on the other hand, the two methods share finitely many approximations and the Classic ones are distinctly more extraordinary, as in the case $k = 43$.

References

[1] S. Mikkilineni. "Continued Fractions and Orbits of a Linear Fractional Transformation." Unpublished pre-print, 2008.

[2] I. Niven, H. S. Zuckerman, and H. L. Montgomery. *An Introduction to the Theory of Numbers.* John Wiley & Sons, 1991, pp. 338–341.

[3] J. Rosen, K. Shankar, and J. Thomas. "Square Roots, Continued Fractions, and the Orbit of 1/0 on $\partial\mathbf{H}^2$." Unpublished pre-print, 2006.