

Continued Fractions and Gaps

Linas Vepstas <linas@linas.org>

12 October 2004 (revised 15 November 2004)

Abstract

Given a continued fraction, we construct a certain function that is discontinuous at every rational number p/q . We call this discontinuity the “gap”. We then try to characterize the gap sizes, and find, to the first order, the size is $1/q^2$, and that, for higher orders, the gap appears to be perfectly ‘randomly’ distributed, in that it is Cauchy-dense on the unit square, and thus, this function has a fractal measure of exactly 2. We find this result to be very intriguing, as we know of no other functions that have this property (There are many fractal *curves* that have this property, but not *functions*. That is, a space-filling curve can be used to enumerate \mathbb{R}^2 by \mathbb{R} but such space-filling curves have locality properties induced by \mathbb{R} that the gap function appears not to have). When examining this function for small rationals, some very curious algebraic relationships appear to relate various rationals.

This paper is part of a set of chapters that explore the relationship between the real numbers, the modular group, and fractals.

1 Continued Fractions and Gaps

1.1 Definitions; Notation

Given a real number $0 \leq x \leq 1$, a sequence of integers $[a_1, a_2, \dots]$ can be found that define its continued-fraction expansion:

$$x = 1/(a_1 + 1/(a_2 + 1/(a_3 + \dots))) \quad (1)$$

Given any particular x , this sequence is straightforward to compute. Rational numbers x always have a finite number of terms in the sequence. For a rational number x , there are two distinct expansions that yield the same value:

$$x_+ = [a_1, a_2, \dots, a_N] \quad (2)$$

and

$$x_- = [a_1, a_2, \dots, a_N - 1, 1] \quad (3)$$

In the following text, we make a simplifying notational confusion: we use the same symbols to denote a sequence and its value when expressed as a continued fraction. Thus, for the above sequences, we may write $x = x_- = x_+$. It will hopefully be obvious which is which.

Given a rational x and an arbitrary complex number w , define new sequences:

$$x_+(w) = [a_1, a_2, \dots, a_N, 1/w] \quad (4)$$

and

$$x_-(w) = [a_1, a_2, \dots, a_N - 1, 1, 1/w] \quad (5)$$

These two functions are analytic in w , almost trivially so, and in fact can be written as $(a + bw)/(c + dw)$ for some integers a, b, c, d . These integers can be readily computed due to an interesting relationship between the Modular Group $SL(2, \mathbb{Z})$ and continued fractions. We give explicit formulas in a different chapter; see also [Cut-the-Knot-UL]. Because of this very simple form, general analytic manipulations are perfectly well-defined on these functions. Clearly,

$$\lim_{w \rightarrow 0} x_-(w) = \lim_{w \rightarrow 0} x_+(w) = x \quad (6)$$

Differentiation w.r.t. w provides a handy tool for constructing various interesting things and proving various limits. We give an exact expression for these derivatives below. Although we've introduced w by appending it to the last term of a rational expansion, we could, if we wanted, introduce w into the N 'th term of an irrational expansion. Most of the statements we make below are for rationals. This is in part because treating the irrationals is a bit harder and confusing, both in these formulas, and in their implementation as algorithms.

There is also an even-odd symmetry for $x(w)$ that we should be aware of, as it plays another important role in the theory. As we point out above, for any given sequence, there is a + and a - generalization. But what if the sequence is already x_- and we choose the - expansion? Then we get

$$x_{--} = [a_1, a_2, \dots, a_N - 1, 0, 1] \quad (7)$$

This expansion invokes a symmetry on $x_-(w)$:

$$x_{--}(w) = [a_1, a_2, \dots, a_N - 1, 0, 1, 1/w] = x_+(w) \quad (8)$$

Thus, we see that a double-expansion is either idempotent or alternating. To be pedantic, we can complete this with the other three expressions: $x_{++}(w) = x_+(w)$ and $x_{-+}(w) = x_{+-}(w) = x_-(w)$. In other words, a double-expansion reflects back, and, in this sense, x_+ and x_- are unique.

If one needs to take differences, e.g. for derivatives, the even-odd expansion is more useful: Define $x_e(w)$ and $x_o(w)$ as follows:

$$x_e(w) = \begin{cases} x_+(w) & \text{if } N \text{ even} \\ x_-(w) & \text{if } N \text{ odd} \end{cases} \quad (9)$$

and

$$x_o(w) = \begin{cases} x_-(w) & \text{if } N \text{ even} \\ x_+(w) & \text{if } N \text{ odd} \end{cases} \quad (10)$$

These two have nice order-preserving properties: if w is a positive real, then

$$x_o(w) < x_e(w) \quad (11)$$

and, for negative real w ,

$$x_o(w) > x_e(w) \quad (12)$$

These two functions also have idempotent or parity-swapping identities:

$$\begin{array}{ll} x_{e+}(w) = x_e(w) & x_{e-}(w) = x_o(w) \\ x_{o+}(w) = x_o(w) & x_{o-}(w) = x_e(w) \\ x_{ee}(w) = x_e(w) & x_{eo}(w) = x_o(w) \\ x_{oe}(w) = x_e(w) & x_{oo}(w) = x_o(w) \\ x_{+e}(w) = x_e(w) & x_{+o}(w) = x_o(w) \\ x_{-e}(w) = x_e(w) & x_{-o}(w) = x_o(w) \end{array} \quad (13)$$

1.2 The Gap

Let us define the *gap* as


$$\Gamma_x(w) = x_e(w) - x_o(w) \quad (14)$$

This gap function is, of course, highly discontinuous in x and analytic in w . It is not hard to discover that for a rational p/q reduced so that p and q are coprime (have no common factors), and $w \leq 1$, the gaps take the form


$$\Gamma_{p/q}(w) = \frac{w}{q^2} \left[2 - w + \frac{w^2}{2} + \frac{w^2}{2} T_{p/q}(w) \right] \quad (15)$$

where the teeth T are bounded: $|T_{p/q}(w)| < 1$ (when $w \leq 1$). For $w = 0$, we have $0 \leq T_{p/q}(0) \leq 1$. A derivation of these results are given in the next section.

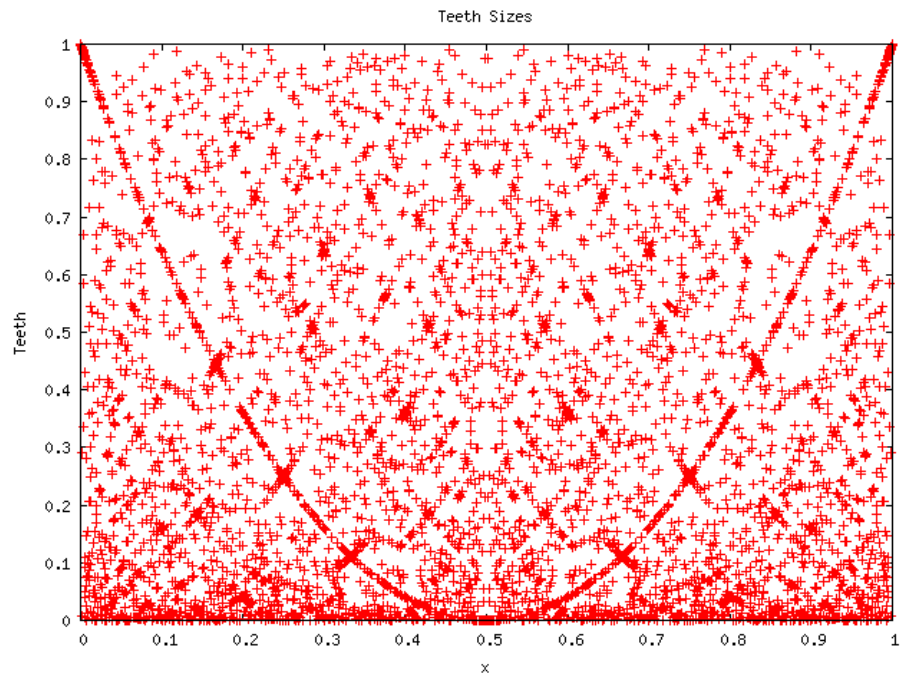
The teeth at first appear to be random: below follows a scatter-plot showing T for a thousand randomly generated small-ish rationals.



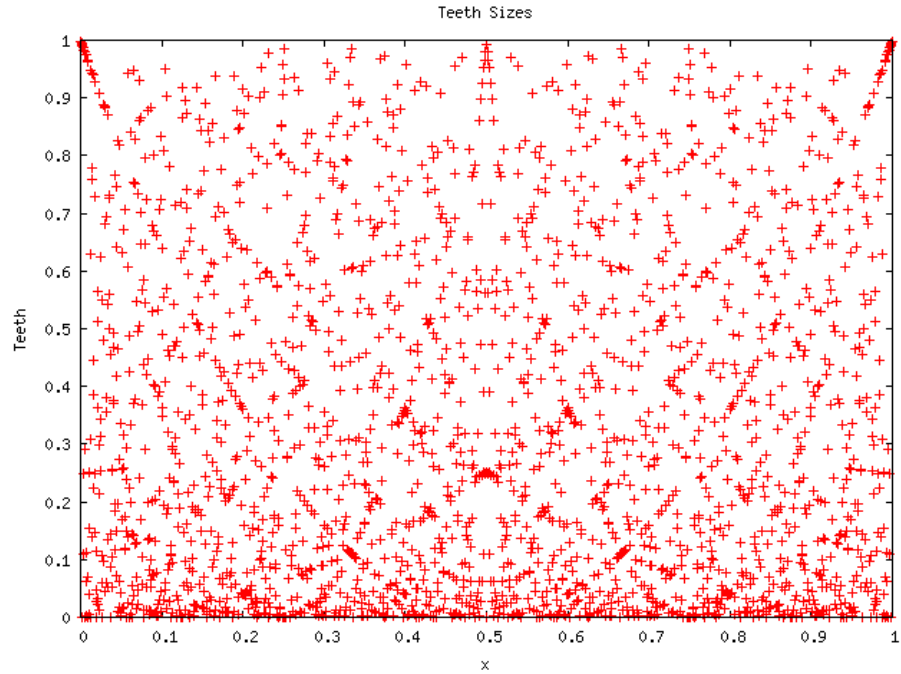
However, closer examination reveals some surprising details. Here's the same picture, but this time showing only the rationals $p/q = [1..719]/720$ (all rationals with numerator between 1 and 719 and denominator of 720):




Note that $720 = 6!$ (six factorial) and that similar but denser images appear for higher factorials: Here's one for $7!$ that is, $p/q = [1..5039]/5040$:



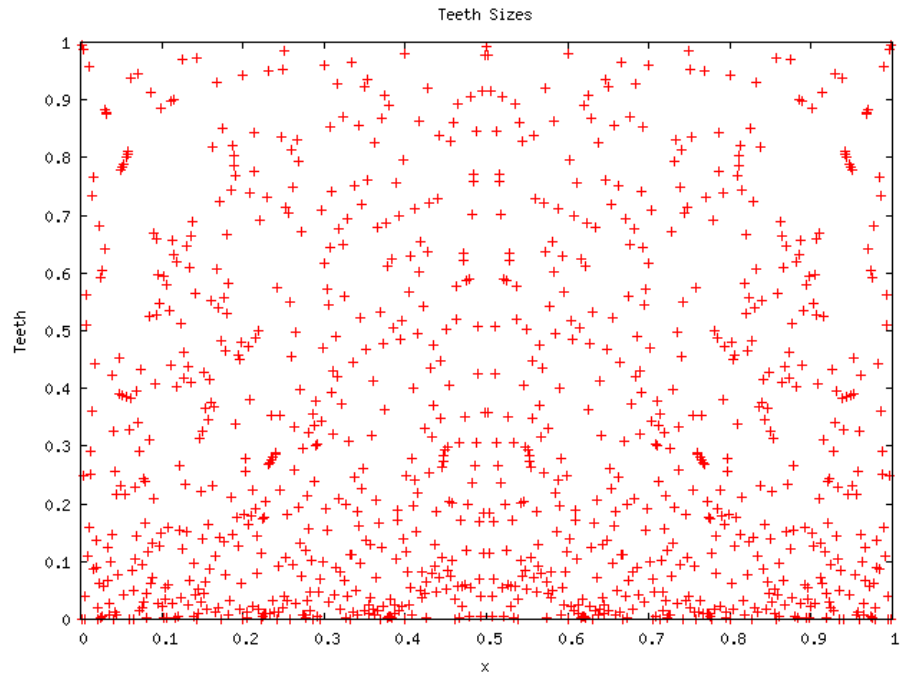
Notice that a Moire pattern of fringes seems to be forming at $x = 1/2$. The appearance of various parabolas at various locations seems to indicate that there are some sort of curious algebraic relationships between the rationals that are worth exploring. The factorial in the denominator seems important: the picture for the denominator $2 \cdot 3 \cdot 5 \cdot 7 \cdot 11 = 2310$ blurs out the features, and mostly looks much more random, as shown below:




Now that we know what sort of pattern to look for, we can find the cross-bars in this picture, for a denominator of 1024:



The analogous picture for a denominator consisting of powers of 3 is less distinct. Pictures for prime denominators appear considerably more random, although containing hints of a different kind of structure. For example, here is a picture with denominator 1023:



The next picture below is a scatter-plot showing all rationals with denominators less than 200. This now clearly exhibits a fantastic self-similar fractal latticework:



We will recognize in later chapters that these curves correspond to hyperbolic maps of the unit interval back onto itself, generated by elements of the modular group $SL(2, \mathbb{Z})$ acting on binary trees. It would be interesting to specify the shapes of all the various bifurcating curves that seem to appear in this latticework, and to give the algebraic equations whose solutions are plotted above. Note that because a curve can seemingly bifurcate into another curve in many places, enumerating all of them would be a trick.

There are more interesting things to be seen, but first we should provide a simpler expression for $T_{p/q}(w)$ at $w = 0$. Recall that for any fixed rational p/q , that $T_{p/q}(w)$ is of the form $(a + bw)/(c + dw)$ for some positive integers a, b, c, d . For any given, fixed p/q , there exists an $\epsilon > 0$ such that there are no poles in $T_{p/q}(w)$ in the disk of radius ϵ around $w = 0$ (Homework: prove this), and thus its analytic on this disk. This justifies an expansion in w ; we'll give an exact expression in a later section, showing that these manipulations are safe. For the continued fraction $x_+(w) = [a_1, a_2, \dots, a_N, 1/w]$ we write the k 'th partial convergent as $r_k(w) = s_k + wt_k + w^2u_k + w^3v_k + \dots$ which obeys the recurrence relation $r_{k-1}(w) = a_{k-1} + 1/r_k(w)$. Substituting, we get the explicit

recurrence relations for each term:

$$s_{k-1} = a_{k-1} + 1/s_k \quad (16)$$

$$t_{k-1} = \frac{-t_k}{s_k^2} \quad (17)$$


$$u_{k-1} = \frac{-u_k}{s_k^2} + \frac{t_k^2}{s_k^3} \quad (18)$$

$$v_{k-1} = \frac{-v_k}{s_k^2} + 2\frac{t_k u_k}{s_k^3} + \frac{-t_k^3}{s_k^4} \quad (19)$$

which are numerically quite tractable. The boundary conditions are given by $r_{+,N}(w) = a_N + w$ for the x_+ expansion, and $r_{-,N}(w) = a_N - w + w^2 - w^3$ for the x_- expansion. We add $a_0 \equiv 0$ to terminate at the other end. Defining even and odd variants as before, we equate the gap

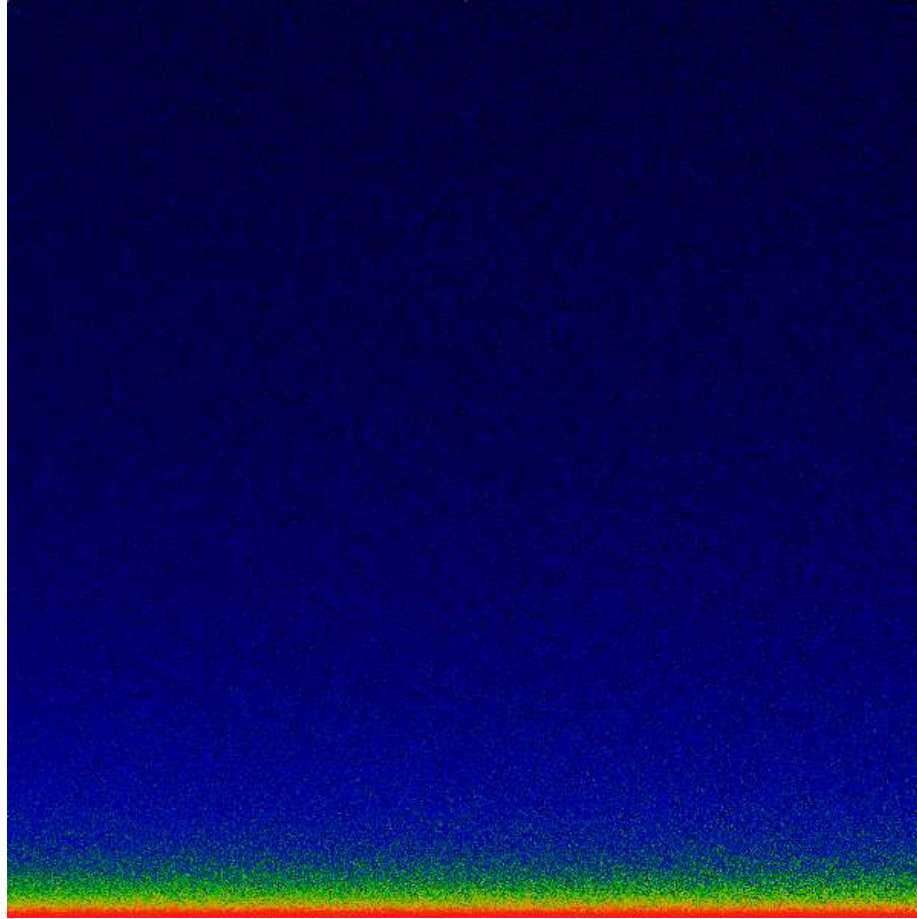
$$\Gamma_{p/q}(w) = w(t_0^e - t_0^o) + w^2(u_0^e - u_0^o) + w^3(v_0^e - v_0^o) \quad (20)$$

Its becomes straightforward to numerically verify that $t_0^e - t_0^o = 2/q^2$ and $u_0^e - u_0^o = -1/q^2$ as we've noted before (Homework: provide a general proof). Far, far more curious is $T_{p/q}(0) = 2q^2(v_0^e - v_0^o) - 1$, of which we drew some pictures above. The pictures below will provide some more surprises. The first is a scatter-plot color-coded as a measure: it shows the distribution of $T_{p/q} \equiv T_{p/q}(0)$ for all rationals with denominator $q \leq 1800$. Black indicates that there were few or no hits to that pixel, then blue, green, yellow and red to denote lots of points hit that pixel. Just as the earlier graphs, shown is a perfect unit square.



For reference, overdrawn over the image is the parabola. The vertical black stripes occur at the Farey numbers, and appear to be ordered in size according to the Farey tree (*i.e.* the largest stripe is at $x = 1/2$, the next largest at $x = 1/3$, $x = 2/3$ and the next smaller ones at $x = 1/4$ and $x = 2/5$ *etc.*). The progression of the horizontal bands are clearly related to the Farey tree through the parabola. Since the parabola gives a rational out when fed one in, we see that we've discovered a different tree of rationals that's related to the Farey tree. One can see that any ratio of polynomials will this give a tree of rationals. We discuss the Farey Tree (or Stern-Brocot Tree) in a different chapter.

The filigree pattern exists only for small rationals. If we start sampling larger rationals, the distribution becomes smooth. The picture below shows a random sampling of rationals with denominators less than 2 million. The image is a bit grainy because the sample size is not large enough; the graininess goes away with larger sample sizes.




This picture tells us something that is surprising: this function is space-filling and dense; it has no holes or islands; it has a fractal measure of 2.0. Even in the land of fractals, this is fairly unusual, as this function is, after all, not some self-similar curve, but rather a 'plain-old' map from the unit interval to the unit interval. By 'dense', I mean the traditional Cauchy-sequence notion of density: for any real values $x, y \in [0, 1]$ and positive $\epsilon > 0$, we can find a rational p/q such that $|x - p/q| < \epsilon$ and $|y - T_{p/q}| < \epsilon$. Of course, a computer generated picture is not a proof, but a proof is straight-forward, and is discussed below. I am not aware of any other map that has this property.

The distribution of the gaps $y = T_{p/q}$ can be easily found to be $1/2\sqrt{y}$ which is sharply peaked at $y = 0$: this is the thin red stripe at the bottom of the above picture. This distribution is due to the Jacobean of a parabola, which we derive below.

1.3 The Gap for Finite w

One can get a better idea of the behavior of $x_\epsilon(w)$ by graphing it, as a function of w , for a number of representative values of x . From the above analysis, we should suspect that for $x = p/q$, the leading w terms are of magnitude $1/q^2$ which would make it hard to

directly compare different values of x . So instead, we introduce a 'normalized' variant: $v_e(x = p/q, w) = x + q^2[x_e(w) - x]$ and similarly $v_o(x, w)$. This is done below, for 200 evenly spaced values of x . On the horizontal axis is x running from 0 to 1, and along the vertical axis, w , from 0 at bottom to 1 at top:



Clearly, it can be seen that $v_e(x, w)$ is a monotonically increasing function of w , which follows easily by examining the partial sums: if $w < z$ then

$$r_N(w) = a_N + w < r_N(z) = a_N + z \quad (21)$$


and

$$r_{N-1}(w) = a_{N-1} + 1/r_N(w) > r_{N-1}(z) = a_{N-1} + 1/r_N(z) \quad (22)$$

Each iteration will reverse the inequality; but v_e by definition has an even number of terms, and thus $v_e(w) < v_e(z)$. QED. Note this would be extremely non-obvious if one just considered that v is a rational polynomial.

One finds a mirror symmetry between the even and odd forms: $v_o(x, w) = -v_e(1 - x, w)$, which is also not quite obvious, given the construction. It is clear from this

picture that the leading w and w^2 terms dominate. We can guess at the form of these terms: we define $\mu_e(x, w) = v_e(x, w) - w + w^2 - w^3/2$ and graph it below, exposing the randomness of the cubic and higher terms directly.



The seeming parallel-ness of some trajectories is an optical illusion, due to the pixelization of the picture. There are more pictures, with a greater number of strands, at the Gap Room <http://www.linas.org/art-gallery/farey/gap-room/gap-room.html>.

1.4 Exact Expressions for the Gap Size

The size of the gap can be given an exact expression in terms of the convergents of the continued fraction. Performing this exercise will help explain some of the fractal behavior seen in the previous sections.

One gets the convergent of a continued fraction by evaluating the fraction only up to the n 'th term, and expressing it as a ratio p_n/q_n . The convergents can be expressed recursively:

$$p_n = a_n p_{n-1} + p_{n-2} \tag{23}$$

and

$$q_n = a_n q_{n-1} + q_{n-2} \quad (24)$$

where we anchor the iteration by defining $p_{-2} = 0$ and $p_{-1} = 1$ and $q_{-2} = 1$ and $q_{-1} = 0$. Iterating, the next few terms are $p_0 = a_0$ and $p_1 = a_0 a_1 + 1$ and $q_0 = 1$ and $q_1 = a_1$. Here we adopted a slightly extended notation from before, defining a_0 as the integer part of the fraction; *i.e.* $x = a_0 + 1/(a_1 + 1/(a_2 + \dots)) \equiv [a_0; a_1, a_2, \dots]$. We then have the convergent:

$$[a_0; a_1, a_2, \dots, a_N] = \frac{p_N}{q_N} \quad (25)$$

Note also that for any $y \in \mathbb{R}$, we have

$$[a_0; a_1, a_2, \dots, a_N, y] = \frac{p_N y + p_{N-1}}{q_N y + q_{N-1}} \quad (26)$$

Note that the convergents have the property $q_n p_{n-1} - p_n q_{n-1} = (-1)^n$. We can use the above to provide exact expressions for $x_+(w)$ and $x_-(w)$. These are:

$$x_+(w) = \frac{p_N}{q_N} + (-1)^N \frac{w}{q_N^2 (1 + w q_{N-1}/q_N)} \quad (27)$$

and

$$x_-(w) = \frac{p_N}{q_N} - (-1)^N \frac{w}{q_N^2 (1 + w (1 - q_{N-1}/q_N))} \quad (28)$$

We can then compute the gap explicitly as

$$\Gamma_x(w) = x_e(w) - x_o(w) = \frac{w}{q_N^2} \left[\frac{2 + w}{1 + w + w^2 \frac{q_{N-1}}{q_N} \left(1 - \frac{q_{N-1}}{q_N}\right)} \right] \quad (29)$$


where no approximation has been made. Expanding the denominator, we get

$$T_{p/q}(0) = 1 - 4 \frac{q_{N-1}}{q_N} \left(1 - \frac{q_{N-1}}{q_N}\right) \quad (30)$$

From this expression, the boundedness of T follows immediately; this comes from the fact that $2q_{N-1} \leq q_N$ (which in turn follows from $2 \leq a_N$ for the last term a_N).


The distribution density $1/2\sqrt{T}$ of the gaps on the unit square can now be understood in terms of the above parabolic formula, as a change-of-variable from the underlying distribution of $R_{p/q} \equiv q_{N-1}/q_N$. Numerically, we can confirm that this distribution is, in a certain sense, perfectly uniform on the half-unit square. Now that we understand that the more fundamental quantity with regard to distributions is $R_{p/q}$, lets repeat some of the earlier graphs. These are shown in 1 and 2. Visually, they don't differ much from thier earlier analogous, except for an effective rescaling of $1/2\sqrt{T}$. Nonetheless, they are presented here for reference.

Figure 1: Ratio of Convergent for Small Rationals




The figure above shows the distribution of the ratio of the final two convergent denominators q_{N-1}/q_N for all rationals with denominators less than 200. Note that the scale along the x-axis runs from 0 to 1 but along the y axis, from 0 to 1/2. To create this distribution, we create an empty grid that is $M \times M = 600 \times 600$ pixels in size. We then consider a fraction $x = p/q$ and its denominator ratio $y = q_{N-1}/q_N$. If we have that $i \leq Mx < i+1$ and $j \leq My < j+1$ for some integers i, j , then we increment the value at pixel (i, j) by one. When we are done, we visualize the grid by assigning black to empty pixels, blue to pixels with a very small count, green to pixels with a medium-small count, yellow to pixels with a medium count, and red to pixels with a large count.

Figure 2: Ratio of Convergents



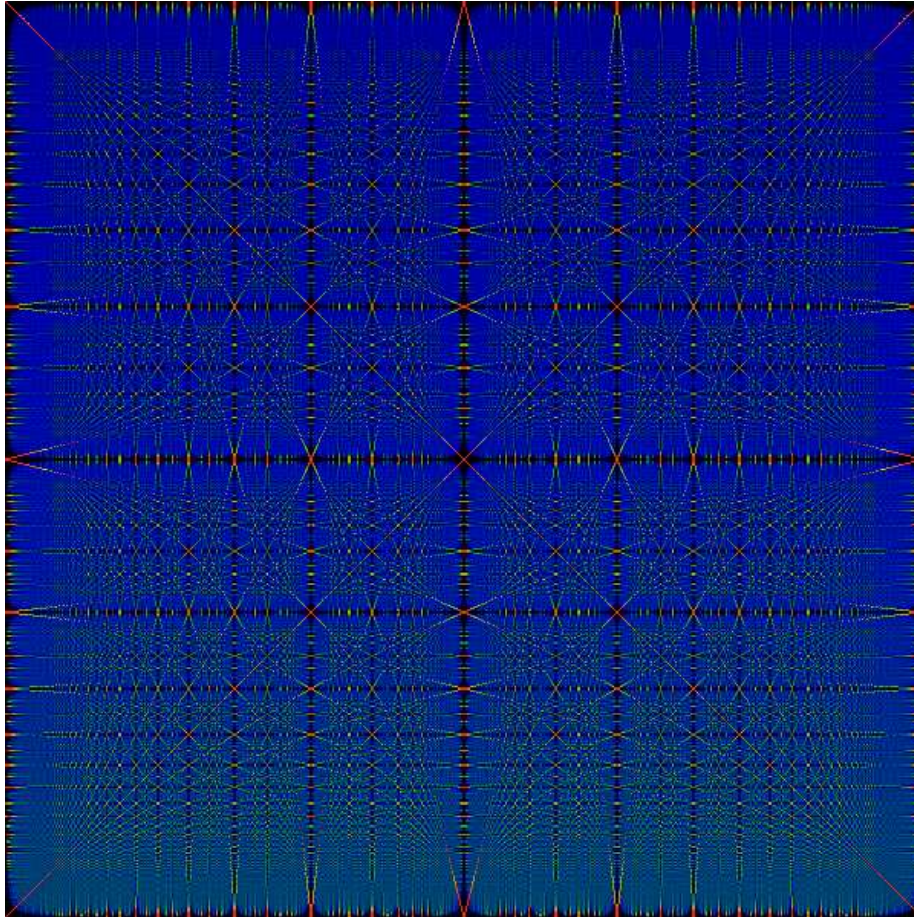
The figure above shows the distribution of the ratio of the final two convergent denominators q_{N-1}/q_N for all rationals with denominators less than 1800. Note that the scale along the x-axis runs from 0 to 1 but along the y axis, from 0 to 1/2. Thus, the largest horizontal line is at $q_{N-1}/q_N = 1/3$, with the remaining lines distributed at the Farey Fractions. The uniform blue color indicates a fairly even distribution, with the coloring as before: red indicating an excess, and black a deficit. As one goes to larger denominators, the uniformity dominates, with the distribution tending towards even-ness.

Figure 3: Ratio of pre-final Convergents




The figure above shows the distribution of the ratio of the pre-final two convergent denominators q_{N-2}/q_{N-1} for all rationals with denominators less than 1800. Note that the scale along the x and y-axis runs from 0 to 1. Note that this figure is fundamentally different, in a certain sense, than the previous figures. Here, as one goes to the higher denominators, one seems to develop both a uniform background, and a filigree superimposed on top. This behaviour is qualitatively different than what one sees when one just considers the final ratios.

Figure 4: High-order pre-final convergents




The figure above shows the distribution of the ratio of the pre-final two convergent denominators q_{N-2}/q_{N-1} for all rationals with denominators less than 16200. The uniform distribution, with a pattern overlay persists to high orders.

Figure 5: Second Convergents




The figure above shows the distribution of the ratio of the two convergent denominators q_1/q_2 for all rationals with denominators less than 1800, and which have at least two terms (so that $q_2 \neq q_1$).

Figure 6: Third Convergents




The figure above shows the distribution of the ratio of the two convergent denominators q_2/q_3 for all rationals with denominators less than 1800, and which have at least three terms (so that $q_3 \neq q_2$).

Figure 7: Fourth Convergents



The figure above shows the distribution of the ratio of the two convergent denominators q_3/q_4 for all rationals with denominators less than 1800, and which have at least four terms (so that $q_4 \neq q_3$).

Figure 8: Fifth Convergents



The figure above shows the distribution of the ratio of the two convergent denominators q_4/q_5 for all rationals with denominators less than 1800, and which have at least five terms (so that $q_5 \neq q_4$).

1.5 Other Convergents

It is of some interest to examine distribution of other convergents. This section reviews some of these.

1.6 Relationship to the Modular Group

The relationship between convergents $q_n p_{n-1} - p_n q_{n-1} = (-1)^n$ implies that the matrix

$$\begin{pmatrix} p_n & p_{n-1} \\ q_n & q_{n-1} \end{pmatrix} \in S^*L(2, \mathbb{Z}) \quad (31)$$

where $S^*L(2, \mathbb{Z})$ is the special linear group of two-by-two matrices over the integers, with determinant equal to plus or minus one. The star on the S is to distinguish it from $SL(2, \mathbb{Z})$, the subgroup having determinant plus one.

1.7 Structure and Randomness

The computer efforts at first seem to paint two contradicting interpretations about the distribution of the gaps. For small denominators, there seems to be a detailed and fine structure; yet, for large denominators, this structure seems to blur away, and one gets the impression of a perfectly uniform distribution. In fact, both conclusions are correct. Each of the pictures above and below were created with a uniform pixel size: 600 pixels dividing up the interval 0 to 1. For small denominators, each pixel may contain only one gap, and maybe none; for large denominators, each pixel may represent the average of dozens or hundreds of gaps. It seems reasonable to conclude that the structure of horizontal and vertical bars is in fact scale-invariant, and persists down to infinitesimal scales. Looking at large-denominator distributions with only 600x600 pixels effectively blurs all the structure away. Thus, on a coarse-grain, the distribution of the gaps appears to be perfectly random; a finer choice of binning, while holding denominators fixed, would lead to pictures of detailed structures at finer levels, and so on.

Thus, in order to talk about the gaps for “all rationals”, one has two competing limits that give different answers. In one case, we take the size of the pixels smaller and smaller, and find, for any fixed scale of denominators, that there is a highly fractal filigree. In the other case, we hold the size of the pixels fixed, and take the limit of larger and larger denominators, and find a perfectly uniform distribution. It seems impossible to define the distribution “in and of itself” without resorting to talking about pixels at some point.

Conjecture: The distribution of the ratio of the partial convergents $R_{p/q} \equiv q_{N-1}/q_N$ is perfectly uniform on the half-unit square, using conventional notions of measure, density and distribution, when we consider all possible rationals, rather than rationals with small denominators.

Proof: None (currently) supplied.

It's not clear how to start on this proof without first developing tools that deal with the scaling relationship of the structures. Blurriness presumably sets in when there are two or three gaps per pixel; thus one will see structure whenever the denominators are at about the same scale as the pixel widths. Thus, the two scale parameters are the pixel size, and the ratio of pixel size to 'typical' denominator magnitude. A proof of uniformity then becomes (presumably) an argument about the mean-square variations of the distribution as a function of the pixel/denom ratio. We expect the mean-square variations to be independent of the pixel size; i.e. scale-independent, and to depend only on this ratio.

1.8 Proof that Gaps are Cauchy-dense on the Unit Square

The incredible 'randomness' of the distribution, as well as its Cauchy-sequence-density on the unit square can now be easily understood. To do this, consider the iteration of the map

$$h(x) = \frac{1}{x} - \left\lfloor \frac{1}{x} \right\rfloor \quad (32)$$

This map has the property that lops off the leading term of the continued fraction: $h([a_1, a_2, \dots]) = [a_2, a_3, \dots]$. Iterating this map clearly gets one further and further into the continued fraction, and it is clear that two points that start arbitrarily close together will have orbits under this map that eventually become uncorrelated. That is, this map clearly has a positive Lyapunov exponent for all irrationals.

Homework: Compute the Lyapunov exponent for this map.

We can gain some intuition by comparing this map to the Bernoulli map

$$b(x) = 2x - [2x] \quad (33)$$

which has the property of lopping off the leading digit of a binary expansion of x . Again, two "random" irrationals that start out arbitrarily close together will eventually have binary expansions that are uncorrelated and completely 'random'. Even though one may know the first N digits of the expansion, one cannot 'predict' the next digit of the expansion; and this is what we mean when we say 'random' or 'uncorrelated'. The only problem with this is that the language used in the last two sentences is completely loaded. If we "know" the number x , then we "know" the N 'th digit in its binary expansion; thus, how can it be "random" and "uncorrelated"? This is a paradox of deterministic chaos that is worth exploring.

Lets try to restate the paradox. We would like to be able to say, that, when considering the entire set of reals on the unit interval, that, when examining the N 'th binary digit in the expansion, that N 'th digit is completely random, and is uniformly distributed (equipartition of probability), with the probability of the digit being 0 being 1/2 and the probability of it being 1 is also 1/2. But of course, this statement is patently false, and therein lies the paradox. We can, of course, trivially and exactly predict the N 'th digit, as the Bernoulli process is completely deterministic. The N 'th digit is

trivially $e_N(x) = \lfloor 2^N x \rfloor - 2 \lfloor 2^{N-1} x \rfloor$ and the graph of $e_N(x)$ is a square-toothed comb. One has this problem even for N getting very large, and approaching infinity: the comb remains uniform and entirely “predictable” no matter how large N gets.

I suppose there are two ways out of this paradox. One way out is to say something like “Yea, Verily, Chooseth two Irrationals whose Binary Expansions are Completely Uncorrelated”, which sounds like an invocation of the Axiom of Choice. Indeed, the set of computable numbers is countable, and is of measure zero on the real number line. In order to compute the N 'th digit of x , one must somehow already 'know' x . Since the set of of unknowable numbers has measure one on the unit interval of reals, one arguably must 'choose' if one is to get a truly 'random' number. Thus, we can anchor the concept of randomness of bit-sequences of the binary expansion of real numbers in the concept of Turing-uncomputable numbers and the Axiom of Choice for choosing one of these unknowable numbers. This is clearly dangerous territory, and somewhat tangled as a definition of randomness.

The other way out of the paradox is far more mundane, and seems constructive, and that is to apply shop-worn statistical methods. After all, chaos occurs in numerical simulations, on finite and computable sets, and not in set-theoretical limits. Chaos is a computational phenomenon. What we can say is that for any given real y and $\delta > 0$ and $\epsilon > 0$, one can always find an integer N such that the average value of e_N approaches $1/2$ on the interval: that is, there exists an integer N such that

$$\left| \frac{1}{2} - \int_y^{y+\delta} e_N(x) dx \right| < \epsilon \quad (34)$$

holds true for all real y and $\delta > 0$ and $\epsilon > 0$. This is a constructive definition of randomness on the real number line that does not require an appeal to the Axiom of Choice or to the choosing of uncomputable (unknowable) numbers as a basis for the randomness of bit-sequences in the binary expansion of a real number. We can approach randomness through traditional delta-epsilon proofs on the expectation values of well-defined quantities. For the following, we are interested in “random” continued fractions, and by analogy, we extend the above definition of “randomness” to the sequence of digits in a continued fraction expansion.

Now, to move the conversation back to the iterated continued fraction map, and that proof.

Theorem: For any real values $x \in [0, 1]$ and $y \in [0, 1/2]$ and positive $\delta, \epsilon > 0$, we can find a rational p/q such that $|x - p/q| < \delta$ and $|y - q_{N-1}/q_N| < \epsilon$.

Proof: (partial sketch) Start by picking $\delta' < \delta$ and some rational p'/q' such that $|x - p'/q'| < \delta'$. Develop the continued fraction expansion of this rational as $p'/q' = [a_1, a_2, \dots, a_{N-2}]$. Then we need to show that we can always pick a positive integer k such that the rational $p''/q'' = [a_1, a_2, \dots, a_{N-2}, a_{N-1}]$ satisfies $|p''/q'' - p'/q'| < \delta - \delta'$ or, equivalently, $|p''/q'' - x| < \delta$ whenever $a_{N-1} \geq k$. Next, we note that if we pick $a_{N-1} > k$ (strictly greater than this time) and any $a_N \in \mathbb{N}$, then $p/q = [a_1, a_2, \dots, a_{N-2}, a_{N-1}, a_N]$ satisfies $|p/q - x| < \delta$. Next, we notice that this freedom to pick a_{N-1} and a_N allows us to jiggle around the value of

q_{N-1}/q_N to satisfy the desired result for many values of y . That is, we use the recurrence relation $q_N = a_N q_{N-1} + q_{N-2}$ to deduce that

$$\frac{q_{N-1}}{q_N} = \frac{1}{a_N} \left(1 - \frac{q_{N-2}}{q_N} \right) \quad (35)$$

and thus

$$\frac{q_{N-1}}{q_N} = \frac{1}{a_N} \left(1 - \frac{1}{a_N a_{N-1} + 1 + a_N q_{N-3}/q_{N-2}} \right) \quad (36)$$

and rearranging,

$$\frac{q_{N-1}}{q_N} = \frac{1}{a_N + 1/(a_{N-1} + r)} \quad (37)$$

In this expression, the value of $q_{N-3}/q_{N-2} \equiv r$ was fixed by our initial choice of p'/q' . However, we are free to pick any $a_{N-1} > k$ and any a_N . Note that by picking $a_N = 2$ and a_{N-1} very large, we can approach the upper limit of the interval, and by picking a_N very large, we can approach the lower limit of the interval. We see that for values of y close to $1/m$ for some positive integer m , we can also approach arbitrarily close, thus completing the proof for these values. However, there are still 'holes' that cannot be approached without jiggering the value of r . For that, we need to adjust our initial pick of p'/q' to get the r that we want. We then find that we have to apply induction, in reverse, to get to there. I believe this completes the proof.

To-Do: The above proof still involves some hand-waving, and thus needs to be tightened up. In particular, there are some values of y that are "hard", requiring the inductive step to be invoked. These seem to correspond to the holes in the holes in the lattice for small denominators. It's not obvious that the inductive step is watertight; although the computer work shows it should be doable.

We will study the iterated map $h(x)$ in great detail in later chapters, where we shall find that it is deeply related to the Modular Group $SL(2, \mathbb{Z})$ and thus to the theory of the symmetry of fractals, and to a variety of fascinating topics in number theory, as well as to the celebrated Riemann Hypothesis.

1.9 Contrast to Space-Filling Curves

The work of Cantor shows that the cardinality of \mathbb{R}^2 is the same as that of \mathbb{R} , namely \aleph_1 . This implies that the points of \mathbb{R}^2 can be enumerated by \mathbb{R} . Space-filling curves such as those of Peano or Hilbert can be used to develop that enumeration. However, these curves have a locality property, in that if two points are close to each other in \mathbb{R} , then the points that they enumerate in \mathbb{R}^2 are also close to each other. This is by construction, of course: the curves of peano or Hilbert are inherently continuous. We can, for example, graph the distance of the Hilbert curve from the the x-axis as

a function of the parameter that takes us along the curve. By construction, this is a continuous curve, and it is **not** space-filling.

By contrast, the gap function is not a curve; by construction, it seems to be discontinuous everywhere.

1.10 The Parabolas in the Distribution

The cause of the principal parabola seen in the pictures is now also clear. Blah Blah Blah write this section too.

The main parabola that is visible is at $T_{p/q}(0) = 4(x - 1/2)^2$, which we can solve for directly to obtain $x = q_{N-1}/q_N$ and $x = 1 - q_{N-1}/q_N$. Recalling that by definition, $x \equiv p/q \equiv p_N/q_N$, we find that the rationals on the main parabola are given by $p_N = q_{N-1}$ and $p_N = q_N - q_{N-1}$. Blah blah blah solve these equations.

Homework: There are other parabolas visible as well, offset on other fractions. Describe these as well. Why are these rationals interesting?

The correct avenue for describing the multitude of the parabolas is to pin down the self-symmetry relationships of the gaps. We already know from the study of continued fractions and Farey Fractions that the relevant symmetry group is the Modular Group $SL(2, \mathbb{Z})$. (See the other papers in this series, duhh). Thus, we should start by reporting on how the gaps transform under the action of the Modular Group elements.

1.11 The Hyperbolic Maps in the Distribution

Exploring the shapes of the hyperbolic maps seen in the pictures requires the development of the theory of these maps as the hyperbolic rotations of a binary tree, through the action of the Modular Group. The development of this theory is done in a later section. We only quote the results xxx here. Blah Blah Blah. What rational numbers show up on which curves? The families of the curves. Write this. So this is another to-do; maybe its own chapter. Again, a deep relationship to the modular group.

1.12 Conclusion

The seemingly pure randomness of the gap sizes is quite intriguing, and suggests possible relationships to similar phenomena in other areas. The (Big and little) Picard theorems comes to mind. The little theorem states that an entire function will attain all possible values, save one or two. The Big Picard theorem states that a function with an essential singularity will attain all possible values (with one or two exceptions), infinitely often, within a finite domain of the singularity. Although the gap isn't analytic, I still find the Picard theorems suggestive. We will develop the tools to analyze the gap in later chapters; in the meanwhile, we note that $h(x)$, or, if you prefer, $\sin(1/x)$, has an essential singularity at zero, and is instrumental in the construction and description of continued fractions.

We also are reminded that cryptographic hashes depend on the ability to distribute points randomly on the unit square, although they do so only on a finite-sized (but large)

lattice. We know that the modular group is central to the study of chaotic dynamical systems and fractals; we find it intriguing that the modular group also underpins the study of elliptic curves, which lead to concepts of elliptic-curve cryptography. Perhaps the seemingly inherent orderly randomness seen in each are different manifestations of the very complex and surprising structure of the modular group.

I'd like to go so far as to propose the 'Fundamental Theorem of Cryptography': Since the gaps are Cauchy-dense on the unit interval, one is always free to pick any sequence of points p/q that one may possibly wish, and then use the $R_{p/q}$ as their cryptographic encoding. That is, the sequence of p/q is the plain-text, whereas the sequence $R_{p/q}$ is the crypt-text. This encoding is provably unbreakable for any plain-text because there exist an infinite number of possible alternate decodings of $R_{p/q}$ that come arbitrarily close to the crypttext, but are arbitrarily distant from the plaintext p/q ; indeed, this is what it means for a set of points to be Cauchy-dense on a plane. What is truly remarkable here is that $R_{p/q}$ is a function, and not some fractal, space-filling curve.

As a final bit of madness, let us note that the gap contains all possible stochastic processes on the unit interval. That is, given a stochastic process, and the usual $\delta, \epsilon > 0$, one can find a sequence of strictly (monotonically) increasing rationals p/q that encode the 'information' in that stochastic sequence to some arbitrarily accurate level. But, by construction, this sequence inherits the modular group symmetry of the continued fractions. Thus, in a certain sense, one might say that stochastic processes have a (perfectly) hidden modular group symmetry. The symmetry is revealed only when one tries to construct things out of the the stochastic sequence; taking, for example, Diffusion Limited Aggregation (DLA). The self-similarity and scaling properties become manifest. In the case of DLA, the dendrites are inherently self-similar. By 'decoding' a pair of dendrites to their representative rationals, one then has a map between the two dendrites, given by the modular group, explicitly exhibiting their self-similarity (to within the δ, ϵ of the decoding). One could thus hand-wavingly say, the perfectly random distribution of the gaps is the "reason" why self-similarity and scaling appears in systems constructed out of random numbers. The self-similarity and scaling are really just a manifestation of the deeper (fractal, modular group) symmetry of the rationals themselves.

References

- [Cut-the-Knot-UL] Alexander Bogomolny's Cut-the-Knot: Continued Fractions on the Stern-Brocot Tree <http://www.cut-the-knot.org/blue/ContinuedFractions.shtml> provides a good elementary introduction to continued fractions.
- [Khi64] Aleksandr Yakovlevich Khinchin, *Continued Fractions* (1964), Dover Publications, Mineola NY.