

# UCT

## Program Size Complexity

KLAUS SUTNER

CARNEGIE MELLON UNIVERSITY

SPRING 2024



**1 Program-Size Complexity**

**2 Computability and Randomness**

Prof. Dr. Alois Wurzelbrunft has discovered a special type of algebraic structure, so-called **Wurzelbrunft algebras**.  $W$ -algebras come in two types, tame or wild. In an heroic intellectual effort, Wurzelbrunft has shown that there are exactly 42  $W$ -algebras and he knows 2 tame ones and 3 wild ones. Unfortunately, it seems quite difficult to figure out which of the others are tame.

Wurzelbrunft wonders whether he should try to prove that tameness of  $W$ -algebras is undecidable.

If only he had taken a course in complexity theory . . .

Any decision problem with finitely many instances is automatically decidable, albeit for entirely the wrong reasons.

To wit, we can hardwire the answers in a lookup table:

$$\begin{array}{cccccc} x_1 & x_2 & x_3 & \dots & x_{n-1} & x_n \\ \hline b_1 & b_2 & b_3 & \dots & b_{n-1} & b_n \end{array}$$

Here  $b_i$  is a bit that encodes the answer for instance  $x_i$ .

The problem is that the correct bit-vector  $b_1, b_2, \dots, b_n$  exists, basta. At least somewhere in set theory la-la land.

Alas, we may not know what it is. We know a decision algorithm exists (really, just a lookup table), but we cannot construct it.

We have talked about this issue before, in the context of showing that

A set of natural numbers is decidable  
iff  
its principal function is computable.

Right-to-left runs into a problem: given a program for the principal function, we cannot decide whether its support is finite. If it is finite, there is a trivial algorithm. Otherwise, there is an algorithm using the gaps in the principal function.

We do not know which of the two algorithms works. But, the algorithm always exists ...

We can push this to absurd levels:

Problem: **Riemann Hypothesis (RH)**  
Instance: A banana.  
Question: Is the Riemann hypothesis true?

If you don't like the banana, replace it by a beer-mug. This problem is easily decidable.

**Algorithm I:** eat the banana, return Yes

**Algorithm II:** eat the banana, return No

One of those two algorithms works. Done.

In the logic literature, the term “decidable” is also used in a different way: one says that some statement  $\Phi$  is *decidable with respect to some logical theory  $T$*  if the theory proves either  $\Phi$  or its negation.

In other words,  $T$  contains enough information to settle the status of  $\Phi$ .

In this sense, the Continuum Hypothesis is not decidable over Zermelo-Fraenkel set theory, even with Choice (Gödel and Cohen).

It is best to avoid this double use, call an assertion **independent** of  $T$  instead.

Pure existence in the standard, non-constructive sense (using set theory) is a bit thin, we would like to have some method to determine the bits, to actually construct the answers.

For example, consider only integer polynomials with at most

100 variables, degree 100, coefficients below 100.

It would be nice to have an algorithm that, given one of these finitely many polynomials, determines whether it has an integral root. An honest algorithm, not a lookup table (which could not be written down in this universe).

There is no hope for this, none whatsoever. We don't even know how to handle degree 4 polynomials with 2 variables. Matiyasevic casts a huge shadow.




One way to tackle this issue is to try to come up with some way to measure to the complexity (in the intuitive sense) of finite bit-vectors.

A priori, computation and complexity theory seem to be utterly useless here, everything is trivial in this framework.

For example, for any string  $x \in 2^n$ , we can build a finite state machine  $\mathcal{A}$  that accepts only this string.

Of course, the machine is essentially just the string itself . . .



Annals of Mathematics, 142 (1995), 443–501

## Modular elliptic curves and Fermat's Last Theorem

By ANDREW WILES\*

For Nada, Clare, Kate and Olivia

*Cubum autem in duos cubos, aut quadratoquadratum in duos quadratoquadratos, et generaliter nullam in infinitum ultra quadratum potestatem in duos ejusdem nominis fas est dividere: cujus rei demonstrationem mirabilem sane detexi. Hanc marginis exiguitas non caperet.*

Pierre de Fermat

### Introduction

An elliptic curve over  $\mathbf{Q}$  is said to be modular if it has a finite covering by a modular curve of the form  $X_0(N)$ . Any such elliptic curve has the property that its Hasse-Weil zeta function has an analytic continuation and satisfies a functional equation of the standard type. If an elliptic curve over  $\mathbf{Q}$  with a given  $j$ -invariant is modular then it is easy to see that all elliptic curves with the same  $j$ -invariant are modular (in which case we say that the  $j$ -invariant is modular). A well-known conjecture which grew out of the work of Shimura and Taniyama in the 1950's and 1960's asserts that every elliptic curve over  $\mathbf{Q}$  is modular. However, it only became widely known through its publication in a paper of Weil in 1967 [We] (as an exercise for the interested reader!), in which, moreover, Weil gave conceptual evidence for the conjecture. Although it had been numerically verified in many cases, prior to the results described in this paper it had only been known that finitely many  $j$ -invariants were modular.

In 1985 Frey made the remarkable observation that this conjecture should imply Fermat's Last Theorem. The precise mechanism relating the two was formulated by Serre as the  $\epsilon$ -conjecture and this was then proved by Ribet in the summer of 1986. Ribet's result only requires one to prove the conjecture for semistable elliptic curves in order to deduce Fermat's Last Theorem.

\*The work on this paper was supported by an NSF grant.

... everybody knows which text contains more information.



A good way to think about this is to try to predict “future bits” in the sequence, assuming there is somehow a natural way to extend it (maybe to an infinite string). Yes, that’s not even ill-defined. Still . . .

- $(01)^\omega$
- concatenate  $01^i$ ,  $i \geq 1$
- binary expansion of  $\sqrt{2}$
- random bits generated by a measuring decay of a radioactive source  
<http://www.fourmilab.ch>.

So the last one is a huge can of worms; it looks like we need physics to do this, pure math and logic are not enough. Kiss ZFC goodbye<sup>†</sup>.

---

<sup>†</sup>Unless you can reconstruct physics in ZFC. Good luck.

How about writing a program that generates the finite string in question?

```
long a[35014], b, c = 35014, d, e, f = 1e4, g, h;

main()
{
    for( ; b=c-=14; h=printf("%04ld",e+d/f) )
        for( e=d%=f; g>--b*2; d/=g )
            d = d*b + f*( h ? a[b] : f/5 ), a[b] = d%--g;
}
```

This program compiles (with a few warnings) and running it produces the first 10000 decimal digits of  $\pi$ .

After removal of all the superfluous white-space this program is only 140 bytes long.

Examples like these strings and the  $\pi$  program naturally lead to the question:

What is the shortest program that generates some given output?

To obtain a clear quantitative answer, we need to fix a programming language and everything else that pertains to compilation and execution.

Then we can speak of the **shortest program** (in length-lex order) that generates some fixed output in  $2^*$ .

**Note:** This is very different from resource based complexity measures (running time or memory requirement; Blum type measures). We are not concerned with the time it takes to execute the program, nor with the memory it might consume during execution.

In the actual theory, one uses universal Turing machines to formalize the notion of a program and its execution. But, as we have seen many times, Turing machines are bit unwieldy, so for intuition it is better to think of

- C programs,
- being compiled on a standard compiler,
- and executed in some standard environment.

Why C? Because it is a no-BS language, close to actual hardware.

So, informally we are interested in the shortest C program that will produce same particular target output. As the  $\pi$  example shows, these programs might be rather weird (in fact, really short programs often are bizarre).

Needless to say, this is just intuition. If we want to prove theorems, we need a real definition.

Consider a universal Turing machine  $\mathcal{U}$ .

For the sake of completeness, suppose  $\mathcal{U}$  uses tape alphabet  $\mathbf{2} = \{0, 1, b\}$  where we think of  $b$  as the blank symbol (so each tape inscription has only finitely many binary digits).


The machine has a single tape for input/work/output.

The machine operates like this: we write a binary string  $p \in \mathbf{2}^*$  on the tape, and place the head right before the first bit of  $p$ .  $\mathcal{U}$  runs and, if it halts, leaves behind a single binary string  $x$  on the tape.

We write

$$\mathcal{U}(p) \simeq x$$





## Definition

For any word  $x \in \mathbf{2}^*$ , denote  $\widehat{x}$  the length-lex minimal program that produces  $x$  on  $\mathcal{U}$ :  $\mathcal{U}(\widehat{x}) \simeq x$ .

The **Kolmogorov-Chaitin complexity** of  $x$  is defined to be the length of the shortest program which generates  $x$ :

$$K(x) = |\widehat{x}| = \min(|p| \mid \mathcal{U}(p) \simeq x)$$

This concept was discovered independently by Solomonov 1960, Kolmogorov 1963 and Chaitin 1965.

Think of  $\widehat{x}$  as the ultimate compressed form of  $x$ , the shortest possible description available (at least in the particular environment  $\mathcal{U}$ ).

Note that we can always hard-wire a table into the program. It follows that  $\hat{x}$  and therefore  $K(x)$  exist, for all  $x$ . Informally, the program looks like

```
print "x1x2...xn"
```

No problem. Moreover, we have a simple bound, there is a constant  $c$  such that for any string  $x$  whatsoever

$$K(x) \leq |x| + c$$

But note that running an arbitrary program  $p$  on  $\mathcal{U}$  may produce no output: the (simulation of the) program may simply fail to halt. Of course, non-halting programs are useless as far as Kolmogorov-Chaitin complexity is concerned.

The claim that  $K(x) \leq |x| + c$  is obvious in the C model.

But remember, we really need to deal with a universal Turing machine.

The program string  $p = \hat{x} \in \mathbf{2}^*$  here could have the form

$$p = ux \in \mathbf{2}^*$$

where  $u$  is the instruction part (“print the following bits”), and  $x$  is the desired output.

So the machine actually only needs to erase  $u$  in this case. This produces a very interesting problem: how does  $\mathcal{U}$  know where  $u$  ends and  $x$  starts? After all, everything is just a bunch of 0s and 1s ...

We could use a simple coding scheme to distinguish between the program part and the data part of  $p$ :

$$p = u_1 0 u_2 0 \dots 0 u_r 1 x_1 x_2 \dots x_n$$

Obviously,  $\mathcal{U}$  could now parse  $p$  just fine, we have a **self-delimiting program**. Alas, this seems to inflate the complexity of the program part by a factor of 2. We'll have more to say about coding issues later.

Note that there are other simple possibilities like  $p = 0^{|u|} 1 u x$ . Here the prefix  $p = 0^{|u|} 1$  delimits the program part  $u$ . Again, we seem to be wasting half our bits.

Also note: we can cheat and hardwire any specific string  $X$  of very high complexity in  $\mathcal{U}$  into a modified environment  $\mathcal{U}'$ .

Let's say

- $\mathcal{U}'$  on input 0 outputs  $X$ .
- $\mathcal{U}'$  on input  $1p$  runs program  $\mathcal{U}(p)$ .
- $\mathcal{U}'$  on input  $0p$  returns no output.

Then  $\mathcal{U}'$  is a perfectly good universal machine that produces good complexity measures, except for  $X$ , which gets the fraudulently low complexity of 1. Similarly we could cheat on a finite collection of strings  $X_1, \dots, X_n$ .

Fortunately, the choice of  $\mathcal{U}$  doesn't matter much. If we pick another machine  $\mathcal{U}'$  and define  $K'$  accordingly, we have

$$K'(x) \leq K(x) + c$$

since  $\mathcal{U}$  can simulate  $\mathcal{U}'$  using some program of constant size. The constant  $c$  depends only on  $\mathcal{U}$  and  $\mathcal{U}'$ .

This is actually the critical constraint in an axiomatic approach to KC complexity: we are looking for machines that cannot be beaten by any other machine, except for a constant factor. Without this robustness our definitions would be essentially useless.


It is even true that the additive offset  $c$  is typically not very large; something like a few thousand.

What we would really like is a **natural universal machine**  $\mathcal{U}$  that just runs the given programs, without any secret tables and other slimy tricks. Think about a real C compiler.

Alas, this notion of “natural” is quite hard to formalize.

One way to avoid cheating, is to insist that  $\mathcal{U}$  be tiny: take the smallest universal machine known (for the given tape alphabet). This will drive up execution time, and the programs will likely be rather cryptic, but that is not really our concern.





A small universal machine like the one above (4 states, 6 tape symbols) would seem to be free of any kind of treachery.



Greg Chaitin has actually implemented such environments  $\mathcal{U}$ .

He uses LISP rather than C, but that's just a technical detail (actually, he has written his LISP interpreters in C).

So in some simple cases one can actually determine precisely how many bits are needed for  $\hat{x}$ .

## Proposition

For any positive integer  $n$ :  $K(n) \leq \log n + c$ .

This is just plain binary expansion: we can write  $n > 0$  in

$$k = \lfloor \log_2 n \rfloor + 1$$

bits using standard binary notation.

But note that for some  $n$  the complexity  $K(n)$  may be much smaller than  $\log n$ . There is a more concise description than the binary expansion.

For example  $n = 2^{2^k}$  or  $n = 2^{2^{2^k}}$  requires far fewer than  $\log n$  bits.

## Exercise

*Construct some other numbers with small Kolmogorov-Chaitin complexity.*

How about duplicating a string? What is  $K(xx)$ ?

In the  $\mathbb{C}$  world, it is clear that we can construct a constant size program that will take as input a program for  $x$  and produce  $xx$  instead. Hence we suspect

$$K(xx) \leq K(x) + O(1).$$

Again, in the Turing machine model this takes a bit of work: we have a program  $p$  that generates  $x$ . To build a program for  $xx$  we could run  $p$ , then copy the output. Alternatively, we could try to run  $p$  twice and put the output right next to the first run. Neither method is trivial, since our tape alphabet is fixed. Just try it.

A very similar argument shows that

$$K(x^{\text{op}}) \leq K(x) + O(1).$$

Concatenation is slightly more complicated: we have to be able to determine the parts of the program for  $xy$  that corresponds to  $\hat{x}$  and  $\hat{y}$ .

$$K(xy) \leq K(x) + K(y) + O(\log \min(K(x), K(y)))$$

Say,  $n = K(x) \leq K(y)$ . Then we write down  $n \hat{x} \hat{y}$  where  $n$  is in binary and self-delimiting.

Here is a slightly counterintuitive fact: we can apply any computable function to  $x$ , and increase its complexity by only a constant.

## Lemma

*Let  $f : 2^* \rightarrow 2^*$  be computable.*

*Then  $K(f(x)) \leq K(x) + O(1)$ .*

*Proof.*

$f$  is computable, hence has a finite description in terms of a Turing machine program  $q$ . Concatenate a self-delimiting version of  $q$  with the program  $\hat{x}$ .

□

The last lemma is a bit hard to swallow, but it's quite correct.

Take your favorite exceedingly-fast-growing recursive function, say, Friedman's mindnumbing  $\alpha$  function<sup>†</sup>. There is an almost trivial algorithm to compute  $\alpha(n)$  for any  $n$ , just a while loop and little bit of wordprocessing.

Then  $\alpha(1) = 3$ ,  $\alpha(2) = 11$ . But  $\alpha(3)$  is a mind-boggling atrocity; just giving a lower bound for this number requires something like the Ackermann function. The kind of monster that only exists in recursion theory, not in any other branch of mathematics.

And yet

$$K(\alpha(3)) \leq \log 3 + \text{a little} = \text{a little}$$

---

<sup>†</sup>See the website for notes on Friedman's function.

## Exercise

*Prove the complexity bound of a concatenation  $xy$  from above.*

## Exercise

*Is it possible to cheat in infinitely many cases? Justify your answer.*

## Exercise

*Use Kolmogorov-Chaitin complexity to show that the language  $L = \{x x^{\text{op}} \mid x \in \mathbf{2}^*\}$  of even length palindromes cannot be accepted by a finite state machine.*



Suppose we have a string  $x = 0^n$ .

In some sense,  $x$  is trivial, but  $K(x)$  may still be high, simply because  $K(n)$  is high: printing 0s is trivial, but we need to know how many.

### Definition

Let  $x, y \in 2^*$ . The **conditional Kolmogorov complexity** of  $x$  given  $y$  is the length of the shortest program  $p$  such that  $\mathcal{U}$  with input  $p$  and  $y$  computes  $x$ .

Notation:  $K(x | y)$ .

Then  $K(0^n | n) = O(1)$ , no matter what  $n$  is.

And  $K(x | \hat{x}) = O(1)$ .

## Lemma

$$K(xy) \leq K(x) + K(y|x) + O(\log \min(K(x), K(y)))$$

*Proof.*

Once we have  $x$ , we can try to exploit it in the computation of  $y$ .

As usual, the log factor in the end comes from the need to separate the shortest programs for  $x$  and  $y$ .

□

Note that we can also do  $K(y) + K(x|y) + \dots$

1 Program-Size Complexity

2 **Computability and Randomness**

$K(x)/|x|$  is the ultimate compression ratio: there is no way we can express  $x$  as anything shorter than  $K(x)$  (at least in general; recall the comment about cheating by hardwiring special strings).

An algorithm that takes as input  $x$  and returns as output  $\hat{x}$  is the dream of anyone trying to improve gzip or bzip2.

Well, almost. In a real compression algorithm, the time/space to compute  $\hat{x}$  and to get back from there to  $x$  is also critically important. In our setting, time and space complexity are being ignored completely.

Alas, there is also the slight problem that neither  $K(x)$  nor  $\hat{x}$  is computable.

As is the case with compression algorithms, we cannot always succeed in producing a shorter string.

## Definition

A string  $x \in 2^*$  is  **$c$ -incompressible** if  $K(x) \geq |x| - c$  where  $c \geq 0$ .  
 $x$  is **incompressible** if it is 0-incompressible.

Hence if  $x$  is  $c$ -incompressible we can only shave off at most  $c$  bits when trying to write  $x$  in a more compact form: an incompressible string is generic, it has no special properties that one could exploit for compression.

The upside is that we can adopt incompressibility as a definition of randomness for a finite string – though it takes a bit of work to verify that this definition really conforms with our intuition. For example, such a string cannot be too biased.

A string  $x \in \mathbf{2}^*$  is **Kolmogorov-random** if  $K(x) \geq |x|$ .

So Kolmogorov-random means 0-incompressible.

## Claim

*There are Kolmogorov-random strings of all lengths.*

This is a straightforward application of pigeon hole.

And, like infinite cardinality arguments, it's a bit disappointing: we would like to understand why some strings fail to be compressible.

Having incompressible/random strings can be very useful in lower bound arguments: there is no way an algorithm could come up with a clever, small data structure that represents these strings.

In general, what can we say about  $c$ -incompressible strings? Here is a striking result whose proof is based on simple counting.

### Lemma

*Let  $S \subseteq 2^*$  be a set of words of cardinality  $n \geq 1$ . For all  $c \geq 0$  there are at least  $n(1 - 2^{-c}) + 1$  many words  $x$  in  $S$  such that*

$$K(x) \geq \log n - c.$$

### Example

Consider  $S = \mathbf{2}^k$  so that  $n = 2^k$ . By the lemma, about half the words of length  $k$  are 1-incompressible.

Also, there is at least one Kolmogorov-random string of length  $k$ .

### Example

Pick some size  $n$  and let  $S = \{0^i \mid 0 \leq i < n\}$ . Specifying  $x \in S$  comes down to specifying the length  $i = |0^i|$ . Writing a program to output the length will often require close to  $\log n$  bits.



This lemma sounds utterly wrong: why not simply put only simple words (of low Kolmogorov-Chaitin complexity) into  $S$ ? There is no restriction on the elements of  $S$ , just its size.

Since we are dealing with strings, there is a natural, easily computable order: length-lex. Hence there is an enumeration of  $S$ :

$$S = w_1, w_2, \dots, w_{n-1}, w_n$$

Given the enumeration, we need only some  $\log n$  bits to specify a particular element. The lemma says that for most elements of  $S$  we cannot get away with much less.

### Exercise

*Try to come up with a few “counterexamples” to the lemma and understand why they fail.*

Proof is by very straightforward counting. Let's ignore floors and ceilings.

The number of programs of length less than  $\log n - c$  is bounded by

$$2^{\log n - c} - 1 = n 2^{-c} - 1.$$

Hence at least

$$n - (n 2^{-c} - 1) = n(1 - 2^{-c}) + 1$$

strings in  $S$  have complexity at least  $\log n - c$ .

□

It gets worse: the argument would not change even if we gave the program  $p$  access to a database  $D \in \mathbf{2}^*$  as in conditional complexity.

This observation is totally amazing: we could concatenate all the words in  $S$  into a single string

$$D = w_1 \dots w_n$$

that is accessible to  $p$  as an oracle.

However, to extract a single string  $w_i$ , we still need some  $\log n$  bits to describe the first and last position of  $w_i$  in  $D$ .

A similar counting argument shows that all sufficiently long strings have large complexity:

### Lemma

*The function  $x \mapsto K(x)$  is unbounded.*

*Actually, even  $x \mapsto \min(K(z) \mid x \leq_{\ell\ell} z)$  is unbounded (and monotonic).*

Here  $x \leq_{\ell\ell} z$  refers to length-lex order.

So even a trivial string  $000 \dots 000$  has high complexity if it's just long enough. Of course, the conditional complexity  $K(0^n \mid n)$  is still small, it's the  $n$  that causes all the problems.

As mentioned, it may happen that  $\mathcal{U}(p)$  is undefined simply because the simulation of program  $p$  never halts. And, since the Halting Problem is undecidable, there is no systematic way of checking:

Problem: **Halting Problem for  $\mathcal{U}$**   
Instance: Some program  $p \in 2^*$ .  
Question: Does  $p$  (when executed on  $\mathcal{U}$ ) halt?

So this is really the same as our old version of Halting on empty tape, just as undecidable as the usual versions.

Let's try to understand intuitively why Kolmogorov-Chaitin complexity must be non-computable.

- Given a string  $x$  of length  $n$ , we would look at all programs  $p_1, \dots, p_N$  of length at most  $n + c$  where  $c$  is the right constant to deal with the “print” statement.
- We run all these programs on  $\mathcal{U}$ , in parallel.
- At least one of them, say,  $p_i$ , must halt on output  $x$ .
- Hence  $K(x) \leq |p_i|$ .

But unfortunately, this is just an upper bound: later on a shorter program  $p_j$  might also output  $x$ , leading to a better bound.

But other programs will still be running; as long as at least one program is still running we only have a computable approximation, but we don't know whether it is the actual value.

## Theorem

*The function  $x \mapsto K(x)$  is not computable.*

*Proof.* Suppose otherwise. Consider the following algorithm  $\mathcal{A}$  with input  $n$ , where the loop is supposed to be in length-lex order.

```
read  $n$   
foreach  $x \in 2^*$  do  
    let  $m = K(x)$   
    if  $n \leq m$  then return  $x$ 
```

Then  $\mathcal{A}$  halts on all inputs  $n$ , and returns the length-lex minimal word  $x$  of Kolmogorov complexity at least  $n$ . But then for some constants  $c$  and  $c'$

$$n \leq K(x) \leq K(n) + c \leq \log n + c',$$

contradiction.

□

Consider the following variant of the Halting set  $H$  (for empty tape), and define the Kolmogorov set  $H'$ , the graph of  $K(\cdot)$ :

$$H = \{e \mid \{e\}() \downarrow\}$$

$$H' = \{x\#n \mid K(x) = n\}$$

### Theorem

*$H$  and  $H'$  are Turing equivalent.*



It is easy to see that  $H'$  is  $H$ -decidable.

Given a string  $x$  of length  $n$ , consider all programs  $p_1, \dots, p_N$  of length at most  $n + c$ .

Use oracle  $H$  to eliminate all non-halting ones.

Run the others to completion and pick out the shortest one that returns  $x$ .

In other words,  $K(x)$  is  $H$ -computable.

This is harder.

We have  $H'$  as oracle, and we need to decide whether Turing machine  $\mathcal{M}_e$  halts on empty tape. Let  $n = |e|$  assuming a binary index  $e$ .

Use oracle  $H'$  to filter out the set of compressible strings of length  $2n$ :

$$S = \{ z \in \mathbf{2}^{2n} \mid K(z) < 2n \}$$

Let  $\tau$  be the time when all the corresponding programs  $\widehat{z}$  halt.

We can compute  $\tau$ : we run all programs of length less than  $2n$  until the appropriate ones have terminated and produced  $S$ .

Here is the key fact:

$\tau$  is large enough to resolve the Halting question for  $\mathcal{M}_e()$ .

**Claim:**  $\mathcal{M}_e() \downarrow$  iff  $\mathcal{M}_{e,\tau}() \downarrow$

Assume otherwise, so  $\mathcal{M}_e() \downarrow$  but  $\mathcal{M}_{e,\tau}() \uparrow$ .

Run  $\mathcal{M}_e$  with a clock to determine  $t > \tau$  such that  $\mathcal{M}_{e,t}() \downarrow$ .

But then we can run all programs of size less than  $2n$  for  $t$  steps and obtain  $S$ , and thus an incompressible string  $z' \in \mathbf{2}^{2n} - S$ .

Alas, our computation of  $z'$  shows that  $K(z') \leq n + c$ , contradiction.

□

If you don't like oracles, we can also represent  $K(x)$  as the limit of a computable function:

$$K(x) = \lim_{\sigma \rightarrow \infty} D(x, \sigma)$$

where  $D(x, \sigma)$  is the length of the shortest program  $p < \sigma$  that generates output  $x$  in at most  $\sigma$  steps,  $\sigma$  otherwise. So  $D$  is even primitive recursive.

Recall our convention about truncated computations returning  $\sigma$  as default, so for small  $\sigma$  we have  $D(x, \sigma) = \sigma$ .

At some point we find the first program  $p < \sigma$  that produces  $x$  in fewer than  $\sigma$  steps and we get the approximation  $K(x) \leq |p|$ . From then on, the function is non-increasing in the second argument. It can drop in value a few times as we find smaller and smaller programs later (that take longer to terminate).

This limit definition produces a  $\Sigma_2$  function:

$$K(x) = y \iff \exists s \forall t (s \leq t \Rightarrow D(x, t) = y)$$

Alas, we cannot compute the threshold  $s$  from  $x$ , otherwise  $K(x)$  would be computable.

Similarly, we cannot compute how often the value of  $D(x, \sigma)$  is going to drop.

At any rate,  $K$  fails to be computable, but it is fairly close, just a limit away.