

CDM

Turing Computability

KLAUS SUTNER

CARNEGIE MELLON UNIVERSITY

FALL 2025



1 Turing Computability

2 The Busy Beaver Problem

The pictures are stunningly beautiful, but we have not really explained what they are a picture of.

More precisely, we have a definition of a Turing machine, but we appeal to intuition when it comes to computations of these machines. As far as intuition is concerned, this is fine, but to prove anything we also need clear, formal definitions.

Given any type of machine \mathcal{M} , to define computations of \mathcal{M} , it is a good idea to think about interrupting a computation.

What information is minimally needed to resume the computation later? This is called a **configuration** or **instantaneous description (ID)**.

Then we explain how \mathcal{M} moves from one configuration to the next, the **one-step** relation. Chaining together single steps by induction, we get a many-step relation that formalizes computations of arbitrary length.

We have to add input/output conventions, and then we're done.

Configurations here consist of:

- the current state
- the current head position
- the current tape inscription

We can think of the head position as an integer, and the tape inscription as a map $\mathbb{Z} \rightarrow \Gamma$. So the space of configurations is

$$Q \times \mathbb{Z} \times (\mathbb{Z} \rightarrow \Gamma)$$

For us only finite tape inscriptions matter, only finitely many cells carry a non-blank symbol.

As it turns out, the following, less obvious, representation for finite inscriptions is often better in applications.

We may safely assume that $Q \cap \Gamma = \emptyset$. Use $Q \cup \Gamma$ as an alphabet and strings of the form $\Gamma^* Q \Gamma^+$ to encode configurations.

Definition

A **configuration** or **instantaneous description (ID)** is a word $y p x$ where $x \in \Gamma^+$, $y \in \Gamma^*$ and $p \in Q$:

$$y_m y_{m-1} \dots y_1 p x_1 x_2 \dots x_n$$

means that the read/write head is positioned at x_1 and the tape inscription is $y_m \dots y_1 x_1 \dots x_n$.

In this representation one can use the theory of finite state machines to explain what it means to perform one step in a computation (a finite state transducer is enough).

The last description really captures one particular type of Turing machine: a

TM with a single, two-way-infinite tape

with finite inscriptions of the form



where the head is at $a = x_1$ and the machine is in state p .

Alas, this is just the tip of the iceberg: it is also useful to consider machines that have multiple tapes (each with a separate read/write head) and/or one-way-infinite tapes. In addition, later we will use separate read-only input tapes, and write-only output-tapes (these are always one-way-infinite).

We may safely assume that $n \geq 1$ since we can always let x_1 be the blank symbol.

It also makes sense to choose n and m to be minimal such that yx captures all the non-blank symbols on the tape. We won't bother to make this part of the definition, though, it really does not matter much.

This is subtly different from the model where a tape inscription is a map $\mathbb{Z} \rightarrow \Gamma$: our approach is coordinate free. Usually that is better, but sometimes it is preferable to keep track of the absolute position on the tape.

E.g., this is used in the proof of the Cook-Levin theorem (hardness of Satisfiability).

Next we need to explain a single step in a computation:

$$y p x \xrightarrow[\mathcal{M}]{1} y' q x'$$

Recall that we assume x to be non-empty and let $\delta(p, x_1) = (q, a, \Delta)$. Then the **next configuration** is defined by

$$\begin{array}{ll} y_m \dots y_1 p x_1 x_2 \dots x_n & \xrightarrow[\mathcal{M}]{1} \\ y_m \dots y_2 q y_1 a x_2 \dots x_n & \Delta = -1 \\ y_m \dots y_1 q a x_2 \dots x_n & \Delta = 0 \\ y_m \dots y_1 a q x_2 \dots x_n & \Delta = +1 \end{array}$$

Strictly speaking, the last definition is wrong: on the one hand, we insist that x is not empty.

On the other hand, the definition can break this condition: think about $\Delta = +1$ and $n = 1$. Similarly, for $\Delta = -1$, we need y_1 to be defined. Of course, there is an easy fix ...

In polite company, this is a non-issue. Alas, any theorem prover would strenuously disagree.

Now we extend the “one-step” relation[†] to multiple steps by induction:

- one step

$$C \mid_{\mathcal{M}}^1 C' \iff \text{as on the last slide}$$

- exactly t steps

$$C \mid_{\mathcal{M}}^t C' \iff \exists D (C \mid_{\mathcal{M}}^{t-1} D \wedge D \mid_{\mathcal{M}}^1 C')$$

- any finite number of steps

$$C \mid_{\mathcal{M}}^* C' \iff \exists t (C \mid_{\mathcal{M}}^t C')$$

[†]For us, the relation is actually a function, but in complexity theory one often consider nondeterministic TMs where there are several possible next configurations.

Lemma

The relation $\mid_{\mathcal{M}}^t$ is primitive recursive, uniformly in t .

Meaning that there is a primitive recursive relation $R \subseteq \mathbb{N}^3$ such that

$$R(t, \langle C \rangle, \langle C' \rangle) \iff C \mid_{\mathcal{M}}^t C'$$

Here $\langle C \rangle$ is supposed to be a code number representing the configuration C . Say, let $\Gamma = [n]$ and $Q = [n+1, m]$ and set

$$\langle C \rangle = \langle y_k, \dots, y_1, q, x_1, \dots, x_\ell \rangle$$

From the perspective of abstract computation, the relation $\mid_{\mathcal{M}}^t$ is easy. But the relation $\mid_{\mathcal{M}}^*$ most emphatically is not.

If there is a p.r. bound on the length of the computations of a TM, then the machine computes a p.r. function: we just iterate the one-step operation an appropriate number of steps. Otherwise we just have to run the machine without any idea if and when it might stop.

So the difference between primitive recursive and Turing computable is just one unbounded search, one existential quantifier. This is the critical difference.

Given any input $x = x_1x_2 \dots x_n \in \Gamma^*$, the **initial configuration** for x is

$$C_x^{\text{init}} = q^{\text{init}} \sqcup x_1x_2 \dots x_n$$

A **halting configuration** is any configuration where the state is q^{halt} and the current symbol is a blank. In this case, we are only interested in the non-blank symbols immediately to the right of the head.

$$C_y^{\text{halt}} = \dots q^{\text{halt}} \sqcup y_1y_2 \dots y_m \sqcup \dots$$

This makes sequential composition of TMs very easy.

Our conventions are fairly natural, but there are lots of alternatives:

- Initial configurations could look like $q^{\text{init}}x_1 \dots x_n$.
- In a halting configuration, we could require the machines to erase the tape except for the output block.
- We are using the blank symbol \sqcup to terminate input and output; we could use a special endmarker like $\#$ instead.

It is tedious but straightforward to check that none of this makes any difference: we get the same clone of computable functions.

Exercise

Explain in detail how to deal with blanks in the input.

Exercise

Come up with a way to associate output with arbitrary halting configurations.

Exercise

Can your machines be simulated by machines conforming to our definitions?

A configuration C is **mortal** if $C \xrightarrow[\mathcal{M}]^* C'$ where C' is halting.

Machine \mathcal{M} **halts on input** x if C_x^{init} is mortal.

We say that $y \in \Gamma^*$ is the **output** of the computation of machine \mathcal{M} on input $x \in \Sigma^*$ if

$$C_x^{\text{init}} \xrightarrow[\mathcal{M}]^* C_y^{\text{halt}}$$

So if we wind up in configuration $abbb q^{\text{halt}} _ abc _ accc$ we consider abc to be the output.

We want to use Turing machines to define computable functions.

Recall from our discussion of primitive recursive functions that we will have to deal with partial functions if we want evaluation to be computable.

Turing machine \mathcal{M} **computes** the partial function $f : \Sigma^* \rightarrow \Sigma^*$ if, for all $x \in \Sigma^*$, we have:

- If f is defined on x , then $C_x^{\text{init}} \xrightarrow[\mathcal{M}]{}^* C_y^{\text{halt}}$ and $f(x) = y$.
- If f is undefined on x , the computation of \mathcal{M} on x does not halt.

Definition

A partial function $f : \Sigma^* \rightharpoonup \Sigma^*$ is **(Turing) computable** if there is a Turing machine \mathcal{M} that computes f .

There is a mountain of evidence that, for any reasonable model of computation \mathfrak{M} , it turns out that \mathfrak{M} -computable is equivalent to Turing computable, so it makes sense to simply say computable, without reference to any particular model.

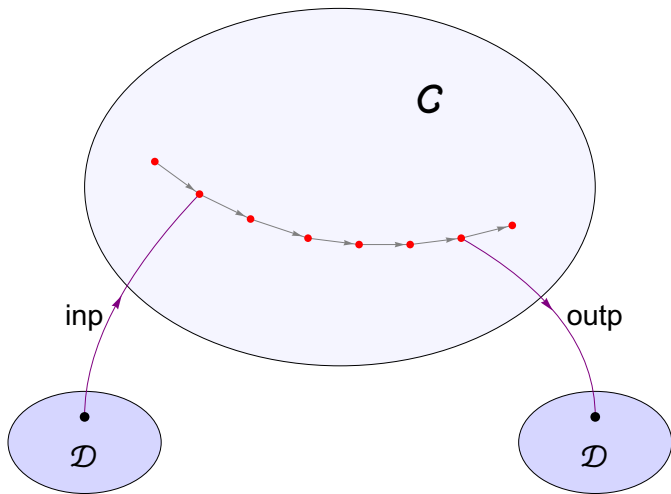
By our definition, Turing computable functions are all unary, they take exactly one string as input.

That's not a problem if we want to compute, say, binary arithmetic functions like addition or multiplication.

$$q^{\text{init}} \sqsubset \#1001\#11000\#\sqsubset$$

We use $\#$ as a separator and write the numbers in binary.

So we will write things like $\mathcal{M}(x, y)$ to indicate TM \mathcal{M} is running on inputs x and y .



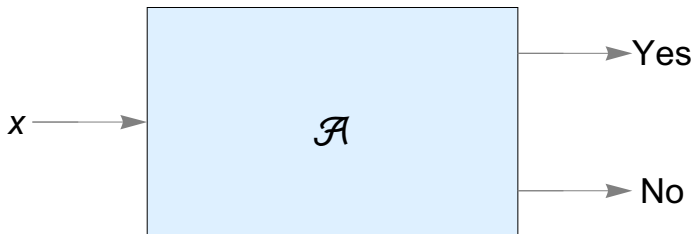
We use characteristic functions to lift Turing computability to relations (or sets).

$$\text{char}_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{otherwise.} \end{cases}$$

Definition

A relation $A \subseteq \Sigma^*$ is **(Turing) decidable** if its characteristic function is Turing computable.

We use the same terminology for $A \subseteq \mathbb{N}^k$ via encoding.



Note that these machines \mathcal{A} always halts.

Informally, a problem is decidable if there is a **decision algorithm** \mathcal{A} that returns Yes or No depending on whether the input has the property in question.

Again, these Turing machines must always halt, and return a one-bit output. They are often called **acceptors**.

By contrast, a Turing machine that computes some arbitrary (and possibly partial function) is called a **transducer**[‡].

[‡]The same terminology is also used for finite state machines.

Lemma

The decidable sets are closed under intersection, union and complement. In other words, the decidable sets form a Boolean algebra.

Proof.

Consider two decidable sets $A, B \subseteq \Sigma^*$. We have two TMs \mathcal{M}_A and \mathcal{M}_B that decide membership.

We can simply run both \mathcal{M}_A and \mathcal{M}_B on input x sequentially, producing a 2-bit result.

It is straightforward to process the output of the two runs and return the correct yes/no answer.



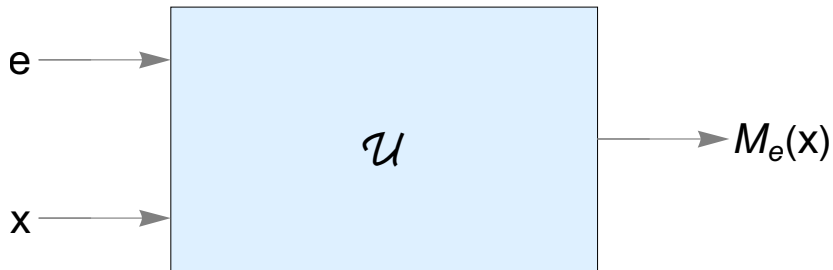
Again there is an **evaluation operation** or a **simulator** that takes as input a Turing machine and a string and runs the machine on the string. The type of eval is

$$\text{eval} : \text{TM}s \times \Sigma^* \longrightarrow \Sigma^*$$

Encoding a TM as a string is trivial. But note, some care is needed to handle the tape alphabet of the machine, we need a little coding to express the arbitrary alphabet in the simulated machine in terms of the fixed alphabet Σ .

At any rate, we can think of a map $\text{eval} : \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$.

So the next question is whether evaluation can be handled by a Turing machine, a so-called **universal Turing machine**.



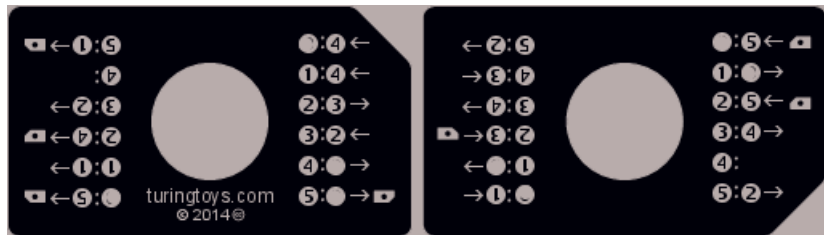
e is the index for the machine and x the input string.

Theorem (Turing 1936)

There is a universal Turing machine.

It immediately follows that the clone of Turing computable functions must contain partial functions.

Finding small universal TMs is still an active research area.



A universal machine (4 states, 6 tape symbols) due to Yuri Roghozin.

Problem: **Halting**

Instance: Index $e \in \Sigma^*$, an argument $x \in \Sigma^*$.

Question: Does Turing machine \mathcal{M}_e halt on input x ?

Problem: **Pure Halting**

Instance: Index $e \in \Sigma^*$.

Question: Does Turing machine \mathcal{M}_e halt on empty tape?

Theorem

The Halting Problem is undecidable.

Proof. Assume a TM \mathcal{H} can decide Halting. Construct the following machine:

```
// machine  $\mathcal{M}$   
// on input  $x$   
  
if  $\text{halts}(x, x)$  // use  $\mathcal{H}$  here  
then halt with output  $0 \frown \mathcal{M}_x(x)$   
else halt with output 0
```

Then \mathcal{M} has some index e and running \mathcal{M} on e produces a contradiction.



Claim: It follows that Pure Halting is also undecidable.

To see why, note that there is a primitive recursive function f such that, given index e and a string x , $e' = f(e, x)$ is the index of a machine that, starting on empty tape,

- first writes x on the tape, and
- simulates machine \mathcal{M}_e on that tape.

It is important there that f is computable, existence of e' alone would not help.

So, we have a notion of a function being computable by a Turing machine that seems to conform very well to our intuition about computability..

This notion does **not change** if we modify our definitions slightly:

- one-way infinite tapes
- multiple tapes
- different head movements
- multiple heads
- different input/output conventions
- different coding conventions

Note that without this kind of robustness our model would be of rather dubious value: each variant would produce a different notion of computability.

We have a perfectly good way to code a Turing machine \mathcal{M} as, say, a sequence number $\langle \mathcal{M} \rangle$.

But the following “encoding” is strictly verboten.

$$\text{code}(\mathcal{M}) = \begin{cases} 2\langle \mathcal{M} \rangle & \text{if } \mathcal{M}(\mathbf{0}) \downarrow \\ 2\langle \mathcal{M} \rangle + 1 & \text{otherwise.} \end{cases}$$

Exercise

Show how to modify our definitions so that Turing machines have total transition functions, but produce the same class of computable functions.

Exercise

Prove that some of the modifications on the last slide similarly yield a type of machine that produces the same class of computable functions as our original Turing machines.

“Prove” here means: think about it for long enough so that you become convinced that an actual proof could be constructed if one really needed a detailed argument.

1 Turing Computability

2 The Busy Beaver Problem

Tibor Radó

On Non-computable Functions

Bell System Technical Journal,
41(3):877-884, 1962.



Radó described a deceptively simple problem in computability. Consider Turing machines on tape alphabet $\Gamma = \{0, 1\}$ and n states.

Question: What is the largest number of 1's any such machine can write on an initially blank tape, and then halt?

Incidentally, it is standard practice to ignore the halting state in the count, so n means “ n ordinary states plus one halting state.” Also, one insists that the tape head always moves.

Rado's original question is somewhat arbitrary, here are two versions more firmly rooted in computability theory.

Time Complexity What is the largest number of moves a halting n -state machine can make?

Space Complexity What is the largest number of tape cells a halting n -state machine can use?

The “number-of-1s” question is below the space complexity, which is below the time complexity.

$BB_H(n)$ is the largest time complexity of any halting n -state machine and refer to BB_H as the **Busy Beaver function**.

$BB_W(n)$ is the largest number of 1's written on the tape when an n -state machine halts.

Clearly, $BB_H(n) \geq BB_W(n)$, but the former has the advantage of relating more directly to the Halting Problem, which one would suspect to be the central issue with busy beaver functions.

$$\text{BB}_H(1) = 1$$

$$\text{BB}_W(1) = 1$$

Make sure to give a precise proof.

$$\text{BB}_H(2) = 6$$

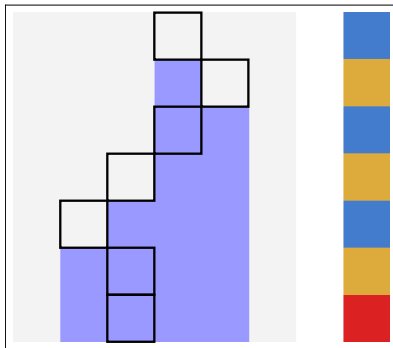
$$\text{BB}_W(2) = 4$$

Amazingly, the answer is no longer obvious.

The champion for both measures is

	0	1
p	q,1,R	q,1,L
q	p,1,L	halt

$p0 \vdash 1q0 \vdash p11 \vdash q011 \vdash p0111 \vdash 1q111$



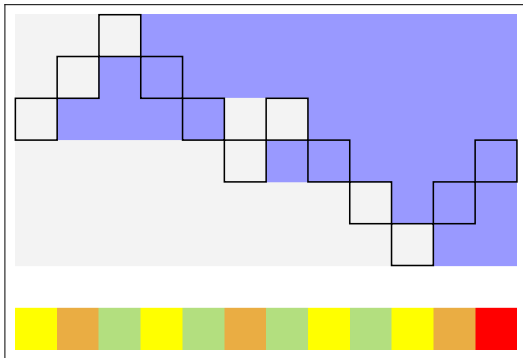
$$\text{BB}_H(3) = 21$$

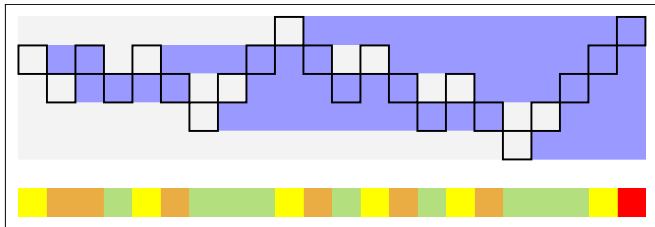
$$\text{BB}_W(3) = 6$$

Here things start to get messy: there are 4 826 809 Turing machines to consider. Exploiting isomorphisms, filtering out machines where all $3 + 1$ states are reachable we get down to 405 072.

Now we just check those for halting and pick out the champion.

Minor Problem: Halting is undecidable and we have to handle each case by hand—somehow concoct an argument that shows that the machine must halt, or that it will not halt forever.





The number of n -state machines grows exponentially:

$$(4n + 1)^{2n}$$

This spins out of control very quickly:

n	#machines
1	25
2	6561
3	4,826,809
4	6,975,757,441
5	16,679,880,978,201
6	59,604,644,775,390,625

These are the raw counts that ignore symmetries between the machines: permuting the states does not produce a “new” machine, nor does switching left and right moves.

Unfortunately, the real problem here is not isomorph-rejection (which requires constructing all machines first), but we need to find a way to only build non-isomorphic ones to begin with.

And, given these numbers, it won't make much of dent no matter what.

Exercise

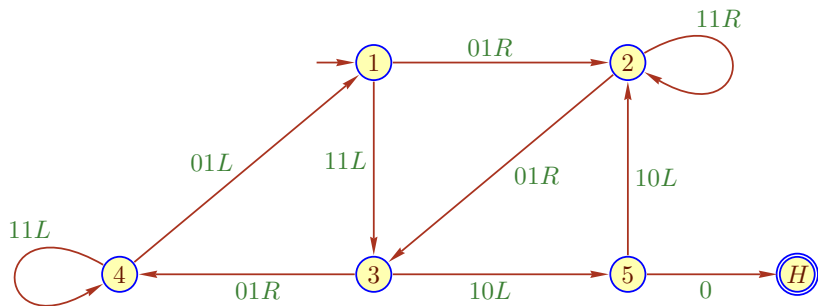
Try to count only non-isomorphic machines.

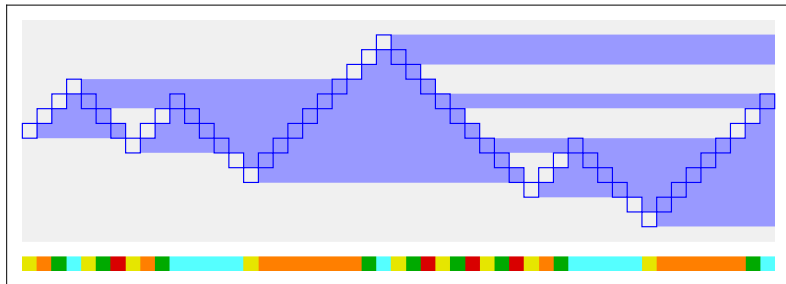
How would you go about generating non-isomorphic machines for $k = 5$?

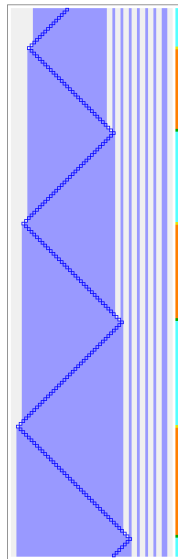
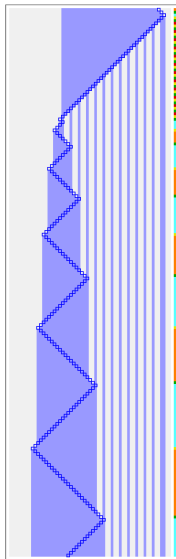
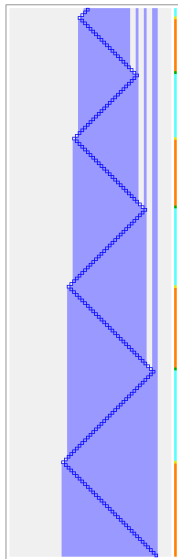
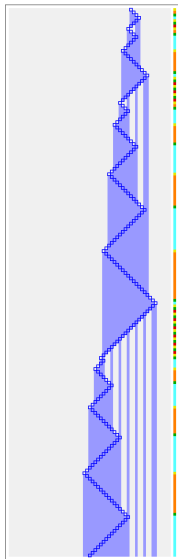
The 5-state champion was found by Marxen and Buntrock in 1989, and its discovery is a small miracle. Here is the table of the machine. Clearly all 5 states plus the halt state are reachable in the diagram.

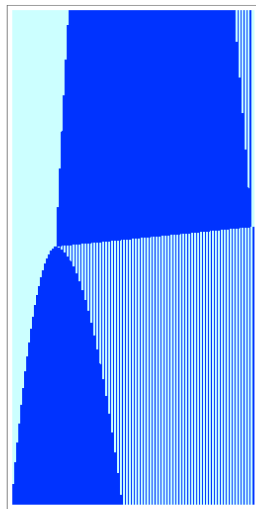
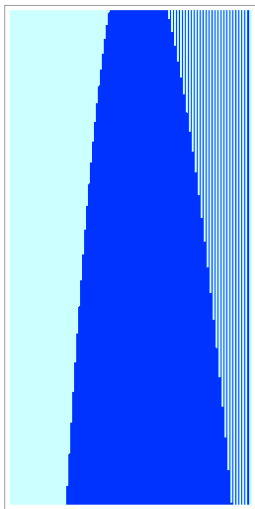
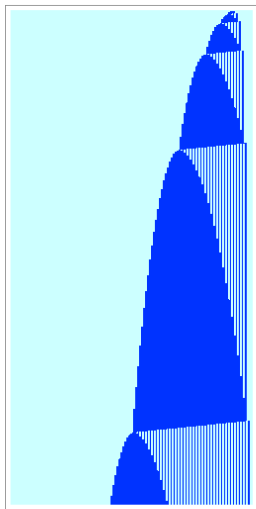
	0	1
1	2,1,R	3,1,L
2	3,1,R	2,1,R
3	4,1,R	5,0,L
4	1,1,L	4,1,L
5	halt	1,0,L

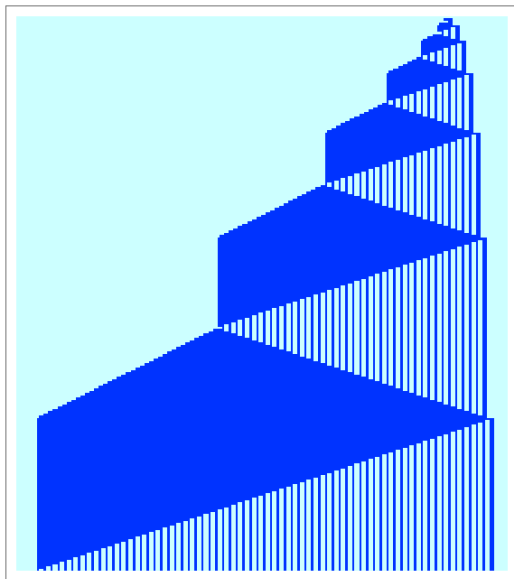
Of course, that's nowhere near enough: they need to appear in the computation on empty tape.











Looking at a run of the Marxen-Buntrock machine for a few thousand steps one invariably becomes convinced that the machine never halts: the machine zig-zags back and forth, sometimes building solid blocks of 1's, sometimes a striped pattern.

Whatever the details, the machine seems to be in a “loop” (not a an easy concept to clarify for Turing machines). Bear in mind: there are only 5 states, there is no obvious method to code the instruction “do some zig-zag move 1 million times, then stop”.

And yet, this machine

stops after 47,176,870 steps

writes $10(100)^{4097}$ on the tape

With a bit of work, one can construe the computation of the MB machine as iterating an arithmetic function

$$f(x) = \begin{cases} (5x + 18)/3 & \text{if } x \bmod 3 = 0 \\ (5x + 22)/3 & \text{if } x \bmod 3 = 1 \\ \uparrow & \text{otherwise.} \end{cases}$$

Starting at 0, the orbit looks like so:

0, 6, 16, 34, 64, 114, 196, 334, 564, 946, 1584, 2646, 4416, 7366, 12284, \uparrow

Big Surprise:

The machine computes these values and then halts.

More precisely, we can filter out configurations of the form

$$C_n = \dots 0 \boxed{1} 0 \underbrace{11 \dots 11}_n 0 \dots$$

n	steps	n	steps
0	15	564	180,307
6	73	946	504,027
16	277	1,584	1,403,967
34	907	2,646	3,906,393
64	2,757	4,416	10,861,903
114	7,957	7,366	30,196,527
196	22,777	12,284	24,576
334	64,407	halt	

For $n = 3k$, the step number is $5k^2 + 19k + 15$,
for $n = 3k + 1$ it is $5k^2 + 25k + 27$.

- BB_H is not computable.
- BB_H has no computable upper bound.
- BB_H dominates all computable functions.

f **dominates** g if $\exists m \forall n \geq m (f(n) > g(n))$.

So busy beaver is all about insanely

There are several fundamental difficulties in computing busy beaver numbers, even for annoyingly small numbers of states.

- Brute-force search quickly becomes infeasible, even for small k .
- The Halting Conundrum:
Even if we could somehow deal with combinatorial explosion, we don't know if a machine will ever halt—it might just run forever.
- Reasoning about the behavior of Turing machines in a formal system like Dedekind-Peano arithmetic or even Zermelo-Fraenkel set theory is necessarily of limited use.

A Π_1 statement of arithmetic ϕ is a formula that can be written in the form

$$\phi \equiv \forall x R(x)$$

where $R(x)$ is primitive recursive[†]. The hard part in checking whether ϕ is true is that there are infinitely many $x \in \mathbb{N}$ to check.

There are quite few open problems in math that are Π_1 .

- The **Goldbach Conjecture** is Π_1 .
- The **Riemann Hypothesis** is also Π_1 .

This is much more surprising; unlike with Goldbach, it requires quite a bit of work to establish this claim.

[†]Written as a formula, it requires only bounded quantifiers and basic arithmetic.

Suppose we have some Π_1 statement $\phi \equiv \forall x R(x)$.

Given a specific natural number n , a Turing machine can easily verify that $R(n)$ holds.

Hence we can build a TM that loops through all values of n and checks $R(n)$. If it ever finds a counterexample, it halts; otherwise it runs forever.

So if we could check Halting, we could handle these Π_1 conjectures.

Allegedly, there is a Turing machine \mathcal{G} on 27 states, binary alphabet, that checks the Goldbach conjecture in this way.

If we could figure out $\text{BB}_H(27)$ we could crack the Goldbach conjecture. In fact, even an upper bound would suffice.

This is a bit of a white lie, $\text{BB}_H(27)$ is a ridiculously large number. In this particular universe we could never even begin to run \mathcal{G} for that many steps—but you get my drift.

Maybe there is a way to avoid all these pesky computational problems?

Instead of computing, we could try to find some clever proof in a sufficiently powerful formal system like Dedekind-Peano arithmetic or Zermelo-Fraenkel set theory[†].

This is most likely wishful thinking. It is currently known (more or less) that $BB_H(n)$ would provide answers to

- $n = 27$: Goldbach
- $n = 744$: Riemann
- $n = 748$: statements independent of ZFC

[†]Of course, a proof is just another computation.

n	$\text{BB}_H(n)$	$\text{BB}_W(n)$
1	1	1
2	6	4
3	21	6
4	107	13
5	47 176 870	4098
6	$> 10 \uparrow\uparrow 15$?

Concrete values are only available for $n \leq 5$, beyond that, we only have bounds. And these bounds are ridiculously large, we have to use special notation to write down the bound even for $n = 6$ in a civilized manner.