Parabola Volume 60, Issue 2 (2024)

The 100 Prisoner Problem

Milan Pahor [1]

1 The Problem

The 100 Prisoner Problem is one of the most bewildering puzzles in the theory of probability. The solution is simple but wildly counter-intuitive. First posed in 2003 by notable computer scientists Anna Gál and Peter Bro Miltersen[2], the problem has been mystifying students and teachers for over twenty years.

A prison warden manages exactly 100 prisoners (wearing individual prison numbers 1 to 100) and decides to collectively give the entire group a chance to win their freedom. In his office, he has a large cupboard with 100 drawers laid out in a 10×10 grid. He randomly distributes the 100 different prison numbers amongst the drawers, exactly one per drawer, and then closes all the drawers.

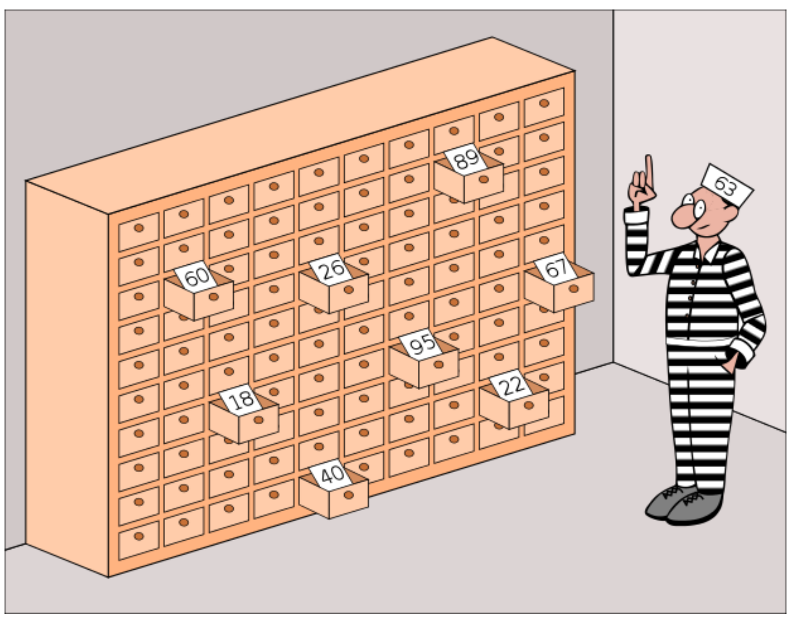

Image from Wikimedia Commons: *100 prisoners problem qtl1.svg*

[1] Milan Pahor lectured in the School of Mathematics and Statistics UNSW Sydney from 1984 to 2020. Following his retirement, he is serving as one of the School's adjunct lecturers.

[2] A. Gál and P.B. Miltersen, The cell probe complexity of succinct data structures, *Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP)*, 2003, pp. 332–344.

Each prisoner, in turn, is given the opportunity to enter the office and locate the drawer which contains their **own** prison number. The catch is that they can only open a maximum of 50 of the 100 drawers. The drawers are all closed again before the next prisoner enters.

If **ALL** of the prisoners locate their respective prison numbers, then the warden will set free the entire population of 100 prisoners! However, if even one prisoner fails to find their own number, then the contest ends immediately and all 100 prisoners will be forced to return to their cells to complete their sentences. It's all or nothing!

The prisoners may discuss the situation before the challenge begins. However, during and after the opening of the drawers, they can no longer communicate with each other. After the process starts, neither the warden nor the prisoners can alter the drawers and their included prison numbers in any way.

Prisoner 27, a crooked accountant, laughs at the warden and says "let's not bother". He explains to his colleagues, quite correctly, that each individual inmate has a probability of success equal to $\frac{1}{2}$. Since there is no communication allowed between the prisoners, the 100 attempts are independent, and thus the probability that they **all** succeed is

$$\left(\frac{1}{2}\right)^{100} \approx 0.000\,000\,000\,000\,000\,000\,000\,000\,000\,000\,79\,.$$

Success is close to impossible!

But Prisoner 63, a Professor in Statistics, claims to have a plan which will increase their chance of freedom from

$$0.000\,000\,000\,000\,000\,000\,000\,000\,000\,000\,79 \text{ to more than } 30\%\,!$$

2 The Professor's Strategy

The Professor outlines her strategy to the rest of the prisoners:

"First we need to number the drawers in some way. I suggest that the upper left drawer be called 1 and then we continue sequentially left to right, so that the top row is 1 to 10, the second row is 11 to 20, continuing down to the bottom row of 91 to 100. We can't put labels on the drawers so this will need to be memorised by us all. The actual order doesn't really matter but we must all use the same drawer numbering.

Now let's consider the plan in my particular case, prisoner number 63. When entering the office I immediately open drawer 63. This drawer will almost certainly not contain my prison number, 63. If it does, then I have succeeded! Suppose instead that drawer 63 contains prison number 18. I then go to drawer 18. If I don't find my own prison number in drawer 18, then I simply move on to the numbered drawer that I *do* find, and so on.

That's it! We simply first open the drawer whose number is our own prison number, and then chase the sequence of numbers, until we either succeed or have exhausted our 50 selections."

Accountant 27 laughs again: “All you are doing is replacing a simple random selection with a more complicated random selection. It’s all still random. The chance that we all find our numbers is still essentially zero”.

But the Professor was right! The prisoners succeeded with their 1 in 3 shot and spent that evening at home with their families. How could this work?

3 The Mathematics

In order to better understand the Professor’s strategy, let’s simplify the situation by reducing the number of prisoners and drawers to 10, with each prisoner allowed to open a maximum of 5 (that is, half) of the drawers.

Mutatis mutandis, the analysis for 100 prisoners will be the same.

Let’s carefully consider one possible scenario:

Drawer Numbers	1	2	3	4	5	6	7	8	9	10
Prison Numbers	4	6	5	7	1	10	3	8	2	9

The table displays the 10 drawers numbered by the prisoners in the top row, and the prison number contents, randomly selected by the warden, in the bottom row. So drawer 3 contains prison number 5; drawer 9 contains prison number 2; and so on.

Note that the numbers in the bottom row are simply a permutation (a re-ordering) of the numbers in the top row. The theory of permutations is very well understood by mathematicians.

If no strategy is used, then each prisoner has a 50% chance of finding their own number. Hence, the probability that a random selection of drawers by the prisoners will result in 10 successes is

$$\left(\frac{1}{2}\right)^{10} \approx 0.001\,.$$

Choosing randomly, the population of 10 prisoners will find freedom roughly one time in a thousand.

But what happens if they follow the Professor’s scheme?

Let’s take prisoner 2 for example. He will go immediately to drawer 2 and be directed to drawer 6, then to 10, then to 9. When he opens drawer 9, he succeeds by finding his own number 2, opening just 4 drawers.

What about prisoner 3? She will start with drawer 3 and be directed to drawer 5, then to 1, then to 4 and then to 7. When she opens drawer 7 she also succeeds by finding her own number 3, opening just 5 drawers.

The strategic journeys of these two prisoners can be displayed in what are called *permutation cycles*. Prisoner 2 is in a 4-cycle and prisoner 3 is in a 5-cycle. Prisoner 8 is in a 1-cycle.

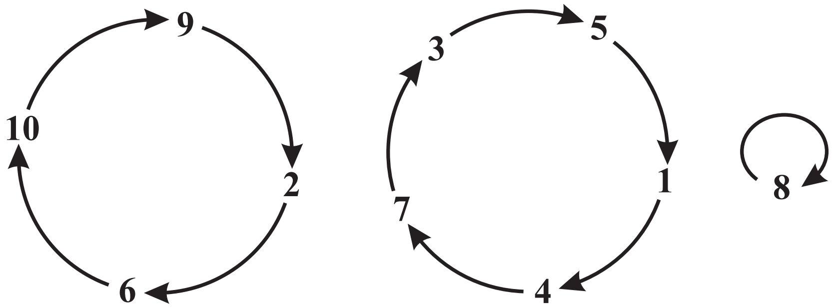

Observe however, that prisoner 9 travels the same cycle as prisoner 2 and will also find their number in 4 turns. Similarly, prisoner 7 shares prisoner 3's cycle and will find his number in 5 turns.

In fact, **ALL** 10 prisoners for this particular permutation will find their prison numbers and, hence, the entire cohort will be freed! Prisoners $2, 6, 10$ and 9 will only need to open 4 of the drawers; prisoners $3, 5, 1, 4$ and 7 will open 5 drawers; and prisoner 8 will find her number immediately.

But will the Professor's strategy always enjoy success? Certainly not! Suppose that the warden's permutation of the prison numbers was as follows:

Drawer Numbers	1	2	3	4	5	6	7	8	9	10
Prison Numbers	5	1	2	10	7	4	9	3	8	6

Then there are only two permutation cycles, a 7-cycle and a 3-cycle.

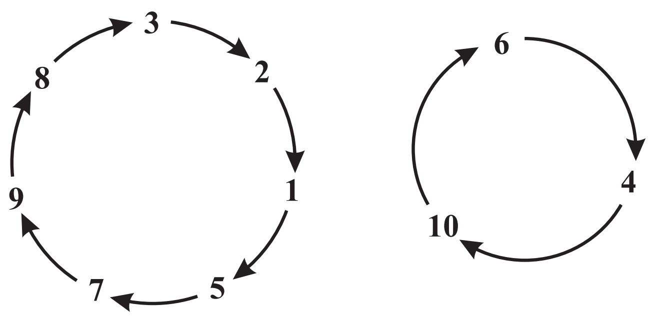

All prisoners in the 7-cycle will fail (in the same way) since they will need 7 attempts to find their own number. For example, the **unsuccessful** journey of prisoner 2 is now

Start at drawer $2 \to 1 \to 5 \to 7 \to 9 \to 8 \to 3 \to$ find prison number 2 .

If 10 inspections are allowed, then each prisoner's trail of drawers will continue, until they close the loop by discovering the drawer which contains their own prison number. This is inevitable, since the number of drawers is finite, and each drawer in the trail (apart from the first) cannot be revisited until the cycle is formed. It is of course true that, when successfully searching for an object, you will always find that object in the last place that you look!

Note in particular that a trail cannot take a short cut and avoid returning to the first drawer by cycling back onto a **part** of itself like this

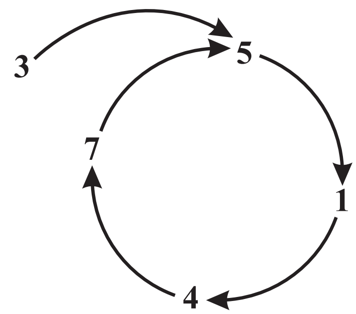

as this would imply that prison number 5 was in **both** drawers 3 and 7.

Each of the 10 drawers contains only one prison number, and each prison number sits in exactly one drawer, so there can be no intersection or overlap between the permutation cycles. **Every** prisoner will be trapped in a single cycle and that cycle will hold their prison number.

The situation is now clear. Every possible permutation of the 10 prison numbers can be uniquely decomposed into a collection of disjoint permutation cycles, and if **ALL** of those cycles are of length 5 or less, then all 10 prisoners will succeed! The prisoners are praying for only short cycles; the warden wants a long cycle. So we really only have one question to consider!

When permuting 10 objects, what is the probability that all permutation cycles are of length at most 5?

Fortunately, we can answer this question using simple permutation and combinations skills! Let's find the probability of a 7-cycle. The other long cycles will behave similarly.

Claim When permuting 10 objects, the probability of a 7-cycle is $\frac{1}{7}$.

Proof Recall that 10 objects can be permuted in 10! different ways. We need to determine how many of these 10! arrangements contain a 7-cycle. First note, that if there is a cycle longer than 5, then there can only be one of them, so there is no need to consider the possibility of multiple long cycles. We first choose the 7 numbers in the cycle. This can be done in 10 choose 7 ways:

$$\binom{10}{7} = \frac{10!}{7!(10-7)!} = \frac{10!}{7!\,3!}\,.$$

Each of these choices of seven may be arranged in a circle in 6! different ways. The remaining three objects can then be arranged in 3! different ways. Thus, there are

$$\left(\frac{10!}{7!\,3!}\right)6!3! = \frac{10!}{7}$$

different permutations of 10 objects which contain a 7-cycle. The probability of a 7-cycle is therefore

$$\frac{\left(\frac{10!}{7}\right)}{10!} = \frac{1}{7}\,,$$

as claimed. QED

The same proof works for all of the long cycles, so the probabilities of 6 ,7 ,8 ,9 and 10 cycles are $\frac{1}{6}$, $\frac{1}{7}$, $\frac{1}{8}$, $\frac{1}{9}$ and $\frac{1}{10}$ respectively.

These events are all mutually exclusive, so the probability of a cycle being longer than 5 is

$$\frac{1}{6}+\frac{1}{7}+\frac{1}{8}+\frac{1}{9}+\frac{1}{10}\approx 0.646\,.$$

Thus, the probability that all the cycles are of length at most 5 is approximately

$$1-0.646=0.354\,.$$

Instead of 1 in 1000, the prisoners have a better than 1 in 3 chance of freedom, if they listen to the Professor!

We can now analyse the cycling strategy for the original 100 Prisoner Problem. In particular, **all** of the prisoners will succeed provided that the maximal cycle length is at most 50. The probability of a cycle being longer than 50 is

$$\frac{1}{51}+\frac{1}{52}+\frac{1}{53}+\cdots+\frac{1}{99}+\frac{1}{100}\approx 0.68817\,.$$

Hence, the probability of freedom for the entire group is approximately

$$1-0.68817=0.31183\,.$$

Amazingly, the Professor's simple, almost trivial strategy, has increased the group's probability of freedom from

$$\left(\frac{1}{2}\right)^{100}=0.000\,000\,000\,000\,000\,000\,000\,000\,000\,000\,79 \quad \text{to} \quad 0.31183\,.$$

Despite the clear maths, this just doesn't seem possible! But it is, and it gets worse: What if we have 1000 prisoners, allowed to make 500 inspections?

The probability of gaining freedom using random selections is then approximately

$$\begin{aligned}0.&000000000000000000000000000000000000000000000000000000000000000000\\&000000000000000000000000000000000000000000000000000000000000000000\\&000000000000000000000000000000000000000000000000000000000000000000\\&000000000000000000000000000000000000000000000000000000000000000000\\&0000000000009332636185032188789900895\,.\end{aligned}$$

This is around the same as the probability of randomly selecting the same atom in the entire universe, five times in a row. Using the Professor's algorithm, the probability is, instead, still close to a third! Indeed it can be shown that, if you start with $2n$ prisoners opening a maximum of n drawers, then the probability of freedom using the cycling algorithm, is always greater than

$$1-\ln(2)\approx 0.307\,.$$

The Professor's strategy mysteriously forces the prisoners into common, cyclic paths of hope, rather than allowing them to flounder in a random scattering of open drawers. Contradicting Nobel Laureate Bob Dylan's words *"When somethin's not right it's wrong"*, the solution to the 100 Prisoner Problem seems to be both right and wrong.

Discussion Question 1

Of course, for the shorter cycles in a permutation of 10 objects, it is not true that the probabilities of 1, 2, 3, 4 and 5 cycles are, respectively,

$$\frac{1}{1}, \quad \frac{1}{2}, \quad \frac{1}{3}, \quad \frac{1}{4} \quad \text{and} \quad \frac{1}{5}.$$

Where does the above argument fail for these cases?

Discussion Question 2

Suppose that the warden has a spy who tells him the drawer numbering system and cycling strategy being used by the prisoners. How could the warden distribute the prison numbers so as to guarantee that **every** prisoner fails?

Discussion Question 3

If the warden is aware of the Professor's strategy, but does **not** know how the prisoners are sequencing the drawers, then what could he do to reduce the prisoner's chance of success?

Discussion Question 4

Suppose that the **first** prisoner is allowed to inspect all 100 drawers and then swap exactly two of the prison numbers as she sees fit. What could she do to guarantee success for the next 99 prisoners?

Discussion Question 5

Given that it is her strategy, Professor 63 is the first to attempt to find her number. She succeeds, but is aware that there is still a two out of three chance of failure for the entire group. She observes the next 49 prisoners also finding their numbers, and then quietly begins packing her bags to go home. How does she know, with half of the prisoners to still take their turn, that success is guaranteed?

Discussion Question 6

If you have some coding skills, then try to run a simulation of the 100 prisoner problem, and see how the probabilities evolve.

Discussion Answer 1

The situation with the shorter cycles is much more complicated due to the lack of uniqueness. For example, if we consider cycles of length 1, then there could be just one of them, or there could be anything up to 10 of them! The probability calculations become mired with inclusion/exclusion arguments and you soon find yourself in the murky world of derangements. The final question in the Australian 2016 H.S.C. Extension 2 examination explores these issues.

Discussion Answer 2

Given this inside information, all the warden would need to do is to put prison number 2 in drawer 1, prison number 3 in drawer 2, all the way down to prison number 100 in drawer 99 and finally prison number 1 in drawer 100. This would doom the prisoners by locking them all in a 100-cycle. Using the Professor's strategy, each prisoner would need to open all 100 drawers to find their prison number.

Discussion Answer 3

The warden would have no options. The situation would then simply be a random permutation with all of the associated probabilistic consequences.

Discussion Answer 4

The first prisoner would simply need to map out all of the cycles. If there are no long cycles, then nothing needs to be done: success is guaranteed. If she finds a cycle longer than 50 (there can only be one), then she only needs to swap two of the prison numbers on opposite sides of the cycle. This will cut the long cycle into two smaller cycles of length less than or equal to 50, once again ensuring success.

Discussion Answer 5

The only way that failure is possible is if a prisoner from the last 50 turns finds themselves in a cycle of length 51 or more. But such a cycle would include at least one prisoner from the first 50, which cannot be, since all of the first 50 succeeded.

If you have any comments or questions regarding this article, feel free to contact me.

Milan Pahor
Adjunct Lecturer
School of Mathematics and Statistics
UNSW Sydney
pahor@maths.unsw.edu.au