Introduction to the general case of the 100 Prisoners Problem

Timothee Schoen

May 17th 2018

Abstract

The 100 prisoners problem is a pretty famous problem in probability theory and combinatorics. It has grown in popularity since 2003, because of its elegant and surprisingly efficient solution to a seemingly impossible riddle. In this paper, we present an introduction to the general case of the 100 prisoners problem. We first introduce for the classic riddle, then we then present a strategy for the general case of the problem and demonstrate the best lower bounds associated with that strategy

Contents

1 The Original Riddle

This problem was first thought of in 2003 by Danish computer scientist Peter Bro Miltersen who published it.

The 100 prisoners problem has different renditions in mathematical literature, but the principle is always the same. Here is the version from *Analytic Combinatorics* (1) :

> The director of a prison offers 100 death row prisoners, who are numbered from 1 to 100, a last chance. A room contains a cupboard with 100 drawers. The director randomly puts one prisoner's number in each closed drawer. The prisoners enter the room, one after another. Each prisoner may open and look into 50 drawers in any order. The drawers are closed again afterwards. If, during this search, every prisoner finds his number in one of the drawers, all prisoners are pardoned. If just one prisoner does not find his number, all prisoners die. Before the first prisoner enters the room, the prisoners may discuss strategy—but may not communicate once the first prisoner enters to look in the drawers. What is the prisoners' best strategy?

2 Solution to the Original Problem

If every prisoners were to choose 50 drawers randomly, then each prisoners would have a probability 0.5 to find their own number and all the prisoners would have a combined probability $(1/2)^{100}$ to be pardoned.

However, there exists a strategy that offers the prisoners a probability of success of more than 30%. The key to that strategy is that each prisoners can use the information from the previous drawers to choose which next drawer to open. The key to this strategy is the probability of success of each prisoner is not independent anymore from the probability of success of the other prisoners.

2.1 Algorithm

The strategy goes as follow:

1. Each prisoner opens the drawer with with its own number.
2. If the prisoner finds its own number in the drawer, he is done.
3. Otherwise, he find an other number in the drawer and he opens the drawer with that number
4. He repeats step 2 and 3 until he is done.

2.2 Analysis of the algorithm

The important part of this special case is that, since every box contains a slip of paper, then each box is going to point to point to exactly one box. By consequence, the boxes can be thought as vertices in a oriented graph, where an edge exists between vertex i and j if only and if the slip of paper in box i points to box j. The distribution of the tickets to the boxes correspond to a random permutation of the numbers from 1 to 100.

Then, we know that the graph will be constituted of only cycles, since every box point to exactly one other box and 2 boxes can not point to the same box. Since the prisoner starts on the box of with his own number, he will be by definition on the cycle that contains his slip of paper. By consequence, he is guaranteed to find his paper after a sufficient number of steps, but he will have to go through the whole cycle before finding his number. By consequence, if the length of the cycle is longer than 50, then the prisoner will fail to find his slip of papers in 50 steps or less.

We now know that the strategy only fails if there is a cycle of length of bigger than 50, because all the prisoners in this cycle will not find their slips of paper in 50 steps or less. Hopefully, since the warden distributed the slip of papers uniformly at random between the boxes (random permutation), then calculating the probability of success is not too hard.

A permutation of the numbers of the numbers from 1 to 100 can only contain at most one cycle of length $l > 50$, where l is the length of the longest cycle. We want to calculate how many different permutations would have a cycle of length l, in order to calculate the probability of having a cycle of that length.

There are $\binom{100}{l}$ ways to select which numbers are in the cycle. Within the cycle, there are $(l-1)!$ ways of organizing the numbers because of the cyclic symmetry. Lastly, the remaining numbers can be arranged in $(100-l)!$ ways. By consequence, the numbers of permutations of the numbers from 1 to 100 with a cycle of length $l > 50$ is

$$\binom{100}{l}(l-1)!(100-l)! = \frac{100!}{l}$$

Since there are a total of 100! possible permutations, then the probability of success of this strategy is

$$1 - \frac{1}{100!}\left(\frac{100!}{51} + ... + \frac{100!}{100}\right) = 1 - \left(\frac{1}{51} + ... + \frac{1}{100}\right) \approx 0.311$$

Even if we increase the number of prisoners to $2n$ and each prisoners can open

n boxes each, the probability of success will still always be lower bounded by $1 - ln2 \approx 0.306$.

2.3 Optimality of the cycle-following strategy

In 2006, Curtin and Warshauer (2) proved the optimiality of the strategy in 2.1 and that the bound in 4.6 could not be improved. Their proof is based on similar game in which the first prisoner opens boxes until he find his slip of paper, then the next prisoner whose slip of paper has not be uncovered open the next boxes and so on. The authors prove that the probability of success in this game is independent of the chosen strategy and is equal to the survival probability in the original problem with the cycle-following strategy. Then, they relate this game to the original riddle through Foata's transition lemma (3) and prove that the cycle-following strategy has to be optimal.

3 The problem of the General Case

This problem was first introduced by Gal and Miltersen (4):

> We have b boxes, labeled 0,1,..., b-1 and a (with $a \leq b$) slips of paper labeled 0,1,..., a-1.The game is played between Player 1 and a team consisting of players $P_0, P_1, ..., P_{a-1}$. Player 1 secretly puts each slip of paper in a different box uniformly at random. Player P_i has to find the slip labeled by at most checking b/k boxes, without communication with any other players (they can still decide on a strategy beforehand, but they can't communicate once they start checking the boxes). The team wins if every player P_i in the team finds the slip labeled with its number.

The 100 prisoners problem is just a specific case of this problem where $a = b$ and $k = 2$. Clearly, we cannot use the same strategy as in part 2.1 if $a \neq b$, because there will be boxes without slips of paper inside.

4 Solution to the General Case

This solution to this general case is inspired by the work of Goyal and Saks (5). The theme of this strategy is to group together boxes in a different bins, then associate each of the a bins to a different slip of paper and finally use a modified version of the algorithm in part 2.1 so that the players can find the slip of paper with their numbers on it.

4.1 Definitions

Definition 4.1. *We are going to imagine that the boxes are organized in a cycle, by consequence, in the rest of this paper, we are going to use the + and − signs as an arithmetic equivalent to + and − mod b.*

Definition 4.2. *For $s, t \in 0, ..., b-1$, we define $[s,t]$ to be the set $\{s, s+1, ..., t\}$ if $s \leq t$ and $[s, b-1] \cup [0,t]$ if $s > t$, as if the boxes were again organized in a cycle.*

Definition 4.3. *We are going to assume that b/a for simplicity purposes and let $d = b/a$. d is the ratio of the number of occupied boxes to the total number of boxes.*

Definition 4.4. *We define $surplus[s,t]$ to be the number of slips in the boxes $[s,t]$ minus $|[s,t]|/d$. It is the difference between the number of occupied boxes in $[s,t]$ and the expected number of expected boxes $[s,t]$.*

Properties:

1. For $s \in [0, b-1]$,
$$surplus[s, s-1] = 0$$
Proof. If we start at box s and go around the full cycle back to s-1, then the surplus will be 0 since the total number of slips is equal to the expected total number of slips.

2. For $s, t, u \in [0, b-1]$,
$$surplus[s,u] = surplus[s,t] + surplus[t+1, u]$$
Proof. It is true because of the additive properties of *surplus*

3. For$i, j \in [0, a-1]$,
$$surplus[di, dj-1] \text{ is an integer}$$
Proof. Since $|[di, dj-1]|$ is a number divisible by d, then $|[di, dj-1]|/d$ and $surplus[di, dj-1]$ will be integers.

4.2 General Strategy

First we partition the set of boxes into a sets called bins: $B_0, ..., B_{a-1}$ of size d where each bin B_i contains boxes $[di, (d+1)i-1]$. The goal will be to associate each one of these bins to a slip of paper.

For $i \in [0, a-1]$, let $m(i)$ be the first integer, such that $surplus[di, m(i)]$ is non negative. The integer $m(i)$ exists as a consequence of Property 1 of surplus.

Since $m(i)$ is the first box where $surplus[di, m(i)]$ is non negative, then we know that the box $m(i)$ has a slip of paper inside, because the surplus changed sign. We define $\pi(i)$ to be the number on the slip of paper found in box $m(i)$.

Lemma 4.5. *π is a permutation of $\{0, ..., a-1\}$*

Proof. By contradiction, let's suppose that there are distinct $i, j \in [0, a-1]$ such that $\pi(i) = \pi(j)$ and by consequence$m(i) = m(j) = m$. We can assume without loss of generality that di precedes dj.

From Proprety 2 of surplus, we have that:

$$surplus[di, m] = surplus[di, dj-1] + surplus[dj, m]$$

$surplus[di, dj-1] \leq -1$, because $surplus[di, dj-1]$ is negative by the definition of $m(i)$ and is an integer by Property 3 of surplus.

Furthermore, $surplus[dj, m] < 1$, because $surplus[dj, m] = surplus[dj, m-1] + surplus[m, m] \leq surplus[m, m] < 1$.

Then, $surplus[di, dj-1] + surplus[dj, m] < 0$ but $surplus[di, m]$ has to be non negative by the definition of m, by consequence there is a contradiction and π is a permutation of $\{0, ..., a-1\}$. This means that every bin is associated to a different slip of paper, we can then use a modified version of the algorithm for the original problem, because we can now find cycles among the bins.

4.3 Algorithm

The algorithm goes as follow:

1. Player P_i starts at the beginning of the bin B_i at box di.
2. P_i sequentially checks the next boxes, keeping tack of the surplus since di until the surplus becomes non negative (at box $m(i)$).
3. P_i opens box $m(i)$, if $\pi(i) = i$, then the player is done.

4. Otherwise, P_i starts again at the beginning of the bin $B_{\pi(i)}$ at box $d(\pi)i)$, resets the surplus and does step 2 and 3 until he finds his number or has opened b/k boxes.

The general idea behind that strategy is close to the special case in part 2.1, where we use the cycles between the slips of papers to get a better solution. This algorithm works, because we use the same principle as in 2.1, where we increase our total probability of success by creating a dependence between the probability of success of each prisoner. However, instead of finding cycles between the boxes like in 2.1, we have to find cycles between bins by sequentially looking for the slip of paper associated with each bin.

4.4 Analysis of the algorithm

Let positive integers a,b,k be the parameters of the game from part 3 such that $b = da$ and $a = 2kn^2$ for some positive integers d and n. Those requirements are not important for the actual analysis, but they help making the proof much cleaner.

Theorem 4.6. *If player 1 randomly assigns the slips of papers to the boxes, then the team of players $P_0, ..., P_{a-1}$ have a probability of winning of at least $2^{-9\sqrt{ka}\log(a-k)}$ if they follow the strategy from 4.3.*

Proof. To find each of the $\pi(i)$, $|[di, m(i)]|$ boxes need to be opened. Let M be the maximum of $|[di, m(i)]|$ for each $i \in [0, a-1]$. The number of boxes opened for any player is upper bounded by M times the length of the longest cycle, because the players are using the cycles in π to find their own slip of paper, like they did in the original riddle.

Since M only depends on the set of occupied boxes (not the actual slips of paper) and the longest cycle in π only depends on the distribution of the slips of paper among the occupied box, then if Player 1 distributed the slip of papers uniformly at random, then M) π are independent. By consequence, we will evaluate independently the probabilities that M and the longest cycle in π are relatively small and multiply those probabilities together to get the lower bound in theorem 4.6.

4.4.1 Bounding the length of the longest cycle

Let $p_1(\alpha)$ be the probability that all cycles in π have length $\leq \alpha$. $p_1(\alpha)$ is trivially at least the probability that all cycles in π have length exactly α (assuming for simplicity that α divides a). Clearly, this new bound is much lower than $p_1(\alpha)$, but it seems difficult to improve it without to many complications.

Lemma 4.7. *The number of random permutation with a/α cycles of length exactly α is:*

$$\frac{a!}{(\frac{a}{\alpha})!\ \alpha^{a/\alpha}}$$

Proof. We choose α elements of the set without ordering a/α times (hence the $a!$), but then since we are working with cycles, we need to divide by α for every subset. Finally we also divide by $(\frac{a}{\alpha})!$ because we do not care about the order of the cycles.

By consequence,

$$p_1(\alpha) \geq \frac{1}{(\frac{a}{\alpha})!\ \alpha^{a/\alpha}} \geq \frac{e^{a/\alpha}}{e\sqrt{2\pi}\ a^{a/\alpha}\ \sqrt{a/\alpha}}$$

The second part of the equation is obtained by using Stirling's approximation of the factorials (6).

4.4.2 Bounding the number of boxes needed to find $\pi(i)$

Let $p_2(\beta)$ be the probability that $M \leq 2d\beta$, with β an integer and $\beta \leq a$. M does not depend on π but only depends on the distribution of the a slips in the b boxes. We now are going to partition the of boxes into a/β sets (we again for simplicity that β divides a) of size $d\beta$ which we are calling groups. We are going to restrict the placement of slips in the boxes such that each group get exactly β slips of paper. This restriction will create a lower bound on $p_2(\beta)$ that could be improved, but it also appears to be difficult to find a better bound.

There are $\binom{d\beta}{\beta}$ possible ways to place β unlabeled slips in the $d\beta$ boxes of a specific group, that we are going to call G for simplicity. We are going to group the different ways to place the slips in G by the maximum of the $surplus[d\beta i, dj-1]$, where $j \in [\beta i, \beta(i+1)]$. In words, it means that we want to find the maximum surplus from the first box in the first bin in the group G, to the last box of a bin in the same group G (we maximize the surplus by choosing which bin) and then group those assignments of slips by the maximum surplus that we found.

Since there are β slips, then the maximum of the surplus of those sets will be at most $\beta - 1$ and at least 0. By consequence, by the the pigeonhole principle, there will be at least $\frac{1}{\beta}\binom{d\beta}{\beta}$ different ways of organizing the slips in G that will be associated to the same maximum surplus. We are going to call this set of the different ways of organizing the slips S, and the maximum surplus reached in this set is going to be called s. We are now going to restrict the placement of slips even more, by forcing the slips in boxes such that the placement is in S (there are at least of such placements $\frac{1}{\beta}\binom{d\beta}{\beta}$).

Lemma 4.8. *If the placement of slips in every groups comes from the set S, then each slip associated with a bin is either in its group or the next one.*

Proof. For a bin B in group G_i, we want to prove that the slip associated with B is either in G_i or G_{i+1}. Let $d\beta i$ and $d\beta(i+1)$ be respectively the index of first box in groups G_i and G_{i+1}. Let j be the index in box in G_{i+1} such that $surplus[d\beta(i+1), j] = s$. We know that such j exists, because G_{i+1} was constructed such that the placement of slips is in S, by consequence, the maximum surplus reached is s.

Let dx be the index of a box B_x located in group G_i. We want to show that $surplus[dx, j] \geq 0$, by consequence the slip associated with bin B_x is located in $[dx, j]$.

By property 2 of surplus,

$$surplus[d\beta i, j] = surplus[d\beta i, dx - 1] + surplus[dx, j] = \\ surplus[d\beta i, d\beta(i+1) - 1] + surplus[d\beta(i+1), j]$$

$$\Rightarrow surplus[dx, j] = surplus[d\beta i, d\beta(i+1) - 1] + surplus[d\beta(i+1), j] \\ - surplus[d\beta i, dx - 1]$$

We know that $surplus[d\beta i, d\beta(i+1) - 1] = 0$ because there are exactly β slips in this interval by construction (because this placement of slips is in S).
We also know that $surplus[d\beta(i+1), j]$ by construction of j.
Finally, $surplus[d\beta i, dx - 1] \leq s$, because the maximum of the surplus is s by construction.
By consequence, since $surplus[dx, j] \geq 0$, then the slip associated with bin B_x is located in the range $[dx, j]$. This lemma is then proved because j is in G_{i+1}.

Since, with those restrictions, the slip associated with a bin is either in its group or the next one, then we know that no bins will required more than $2\beta d$ boxes to be opened to find the associated slip of paper ($|[di, m(i)]| \leq 2\beta d$). Now we want to find the number of assignments of slips that fit with our restrictions. We already know that for each of the a/β groups, they are at least $\frac{1}{\beta}\binom{d\beta}{\beta}$ placements for each group that are in S. Since we have a/β groups, we have a total of

$$\frac{1}{\beta}\binom{d\beta}{\beta}^{a/\beta}$$

possible assignments that fit our restrictions.

Since they are a total of $\binom{da}{a}$ possible assignments without restriction, because

there are $b = da$ boxes and a slips of papers.

$$p_2(\beta) \geq \frac{\frac{1}{\beta}\binom{d\beta}{\beta}^{a/\beta}}{\binom{d\beta}{\beta}} \geq \frac{a^{1/2}}{\beta^{3a/2\beta}}\left(\frac{d}{2\pi(d-1)}\right)^{\frac{a}{2\beta}-\frac{1}{2}} e^{-\frac{a}{6\beta^2}}$$

The right side of the equation is obtained by using Stirling's approximation of the factorials (6) and with a lot of algebra.

4.4.3 Putting the bounds together

We know that the biggest cycle of π has length at most α with probability at least $p_1(\alpha)$. We also know that no bins will required more than $2\beta d$ boxes to be opened to find the associated slip of paper with probability $p_2(\beta)$. By consequence, each player needs to open at most $2d\alpha\beta$ with probability at least $p_1(\alpha)p_2(\beta)$ (since the probabilities are independent). If $2d\alpha\beta \leq b/k$, then team wins with probability $\geq p_1(\alpha)p_2(\beta)$. We assume that $d = b/a$ in part 4.4. We then choose $\alpha = \sqrt{a/2k}$ and $\alpha = \sqrt{a/2k}$ (for optimality of the bound), which are integers as we also assumed that $a = 2kn^2$ in part 4.4.

$$p_1(\sqrt{a/2k})p_2(\sqrt{a/2k}) \geq 2^{-9\sqrt{ka}\log(a-k)}$$

This completes the proof.

The right part of the equation can be obtained with a few steps of algebra that you can find in Goyal's and Saks's paper (5).

5 Discussion

In this paper, we tried to expose the readers to an introduction to the general case of the 100 prisoners problem. Even with the same strategy that we used in part 4.3, we could find some better bonds by tweaking the calculation as Y. Wang did in his paper (7) by finding a new lower bound of $2^{-3\sqrt{ak}\log(a)-\frac{2\log(e)}{3}k}$ on the probability of success.

We might also want to consider other strategies. For instance, it might more natural for player B_i to start at box di and continue looking at the boxes until he finds a slip of paper. But with this strategy, some slip will be hit and the team will by consequence always lose.

6 Acknowledgement

I would like to thank Professor Postinov for allowing me the opportunity to learn more about this really interesting problem. I would like also like to thank the rest of the student in 18.204 for introducing to some great new subjects in mathematics and for their feedbacks on my presentations.

References

[1] R. S. P F.lajolet, *Analytic Combinatorics*. Cambridge University Press, 2009.

[2] M. W. E. Curtin, "The locker puzzle," *Mathematical Intelligencer*, pp. 28–21, 2006.

[3] Wikipedia, "Foata's transition lemma."

[4] A. G. P. B. Miltersen, "The cell probe complexity of succinct data structures," *Proc 30th IntColl Automata, Languages, and Programming (ICALP)*, pp. 332–344, 2003.

[5] M. N. Goyal, "A parallel search game," *Random Structures and Algorithms*, vol. Volume 27, 2005.

[6] Wikipedia, "Stirling's approximation."

[7] Y. Wang, "The locker puzzle," 2015.