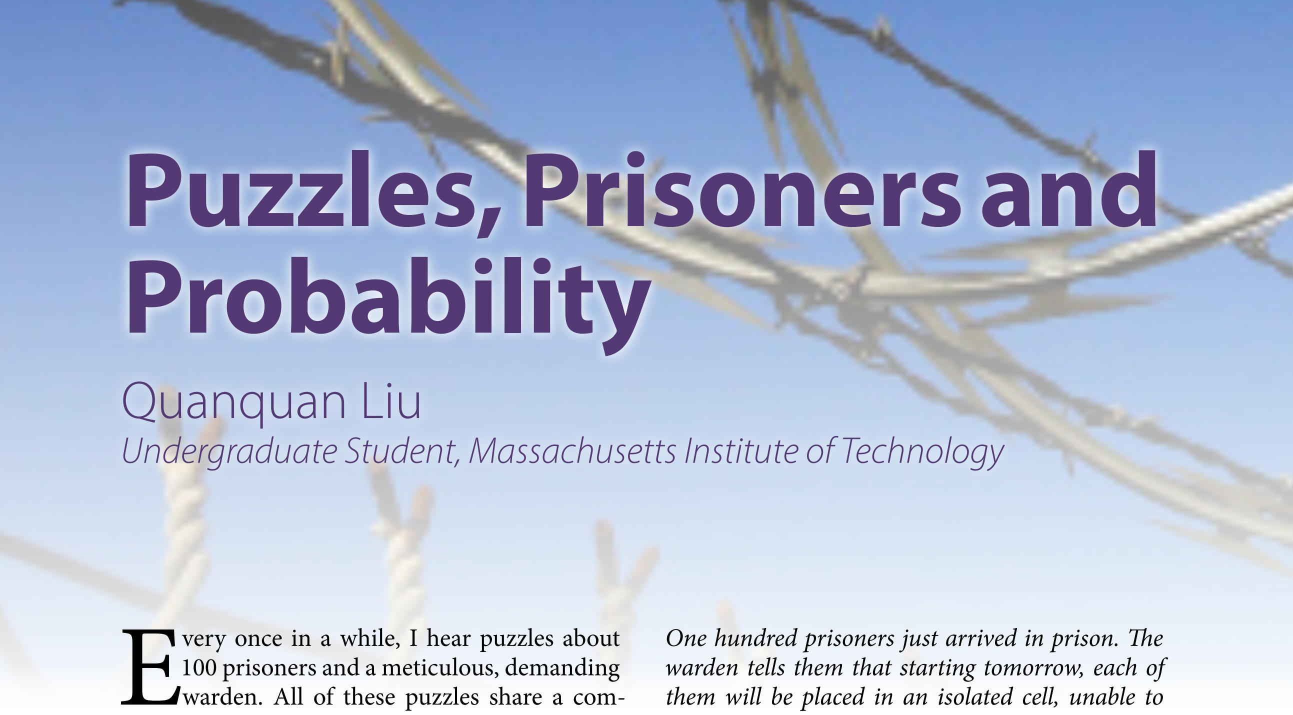

Puzzles, Prisoners and Probability

Quanquan Liu
Undergraduate Student, Massachusetts Institute of Technology

Every once in a while, I hear puzzles about 100 prisoners and a meticulous, demanding warden. All of these puzzles share a common characteristic: a set of prisoners must work together to devise a clever scheme to thwart the warden. I hear these puzzles often enough that each time they reappear, I view them with an increased level of understanding corresponding to the stage of my mathematics education.

Any problem may have a solution, but, sometimes, that solution may not be the most efficient one possible. For example, suppose you want to find something in your room but don't remember where you put it. You can either search the room by yourself. Or you can call your (many) friends to help you. From a correctness standpoint, both solutions are correct. You'll find what you're looking for eventually. But from an algorithmic standpoint, the second solution where you search in parallel with your friends is better because it has a shorter runtime. The same holds for solutions to the 100 prisoners puzzle. Some solutions may be theoretically correct answers to the puzzle but may have expected runtimes that exceed the lifespan of an average person and are, thus, practically undesirable. Now, through this lens, the lens of a theoretical computer science student, I would like to present to you the 100 prisoners puzzle and its variants.

100 Prisoners and a Light Bulb

The original, very famous puzzle involving an interrogation room, a light bulb, and 100 prisoners is the following (paraphrased from Wu in [3]):

One hundred prisoners just arrived in prison. The warden tells them that starting tomorrow, each of them will be placed in an isolated cell, unable to communicate amongst themselves. Each cell has a window so the prisoners will be able to count the days. Each day, the warden will choose one of the prisoners uniformly at random with replacement, and place him in a central interrogation room containing only a light bulb with a toggle switch. The light bulb is initially switched off. The prisoner may observe the current state of the light bulb. If he wishes, he may toggle the light bulb. He also has the option of announcing that he believes all prisoners have visited the interrogation room at some point in time. If this announcement is true, then all prisoners are set free, but if it is false, all prisoners are executed. The warden leaves, and the prisoners huddle together to discuss their fate. Can they agree on a strategy that will guarantee their freedom? [3]

One common solution to the puzzle is to divide the days into 100-day blocks and instruct any prisoner to toggle the light off if he is interrogated twice within the same block. The first prisoner of each block turns the light on and the last prisoner checks whether the light is still on when he enters the interrogation room at the end of the 100-day block. If the light is still on and he did not enter the room on any previous day within the block, he declares that all prisoners have visited the interrogation room [3]. This solution is technically correct because it guarantees the prisoners their freedom, but the prisoners are expected to be freed after 1.072×10^{44} days, in years $\approx 10^{31}$ times the age of the universe. From an algorithmic standpoint, this solution is rather poor because it has an expected runtime of $O(n^{1/2}e^n)$ [3] for n prisoners.

The challenge now is to find a solution that is correct and also has an optimal runtime. Such a solution is more likely to guarantee that the prisoners are freed while they are still alive.

The canonical (better) solution is to designate a "leader" to be the person who counts the number of unique prisoners who have been interrogated. The leader may do so by counting the number of times the light bulb has been switched on. A prisoner who has not yet toggled the light switch will turn the light on if it is currently off. A prisoner will do nothing if he enters the room when the light is currently on. The leader turns the light off each time she leaves the room and increases her counter when she sees a light that is on. Thus, after counting 99, the leader may declare that all the prisoners have been interrogated at least once. (It is sufficient to count to 99 because the leader herself counts as the last prisoner.)

How long are the prisoners expected to wait? Suppose that T represents a counter for the number of times the bulb has been switched on. We may count the expected number of days until $T = 99$. Let X_i denote the number of days that pass between an increment of the counter from when $T = i$ until $T = i + 1$. Let Y_i denote the number of days from when a leader turns off a light bulb until a prisoner turns on the light. Let Z_i denote the number of days from when a prisoner turns on the light until the leader enters the room to see the newly turned on light bulb. Thus, $X_i=Y_i+Z_i$. Let X be the number of days the strategy requires in total before the prisoners are freed. Given n prisoners, the probability of turning on the i^{th} light is $\frac{n-i}{n}$. The probability that the leader enters the room on any day is $\frac{1}{n}$. By linearity of expectation,

$$E[X] = \sum_{i=1}^{n-1} E[X_i] = \sum_{i=1}^{n-1}(E[Y_i] + E[Z_i])$$
$$= \sum_{i=1}^{n-1}\left(\frac{n}{n-i} + n\right) = n^2 - n + nH_{n-1}$$

In asymptotic notation, the "leader" algorithm has an expected runtime of $O(n^2)$ days [3]. When there are 100 prisoners, the expected wait time is 10417.7 days or approximately 29 years [3]. Though still a long time, it is within the prisoners' lifespans.

Wu [3] further summarized some strategies that may lead to even shorter wait times. One such strategy achieves an expected runtime of $O(n(\log n)^2)$. The key insight behind this algorithm is to allow "assistant" leaders to help the leader by doing some of the counting. Then, the leader would sum together the totals of all the "assistant" counts to determine if all prisoners have visited the interrogation room. To do this, we must be able to divide up the counting of the light bulbs into blocks of days. There must be a block for assistants to count the number of prisoners and a different stage for assistants to tell the leader their total [3]. See [3] for more details.

But can we achieve a solution with an even better runtime, for example, a solution with an $O(n)$ expected runtime? Turns out, the answer is no for the $O(n)$ runtime solution. It is a common joke among CS theoreticians that we hate lower bounds because it prevents us from making better algorithms. The reason why we can't create an $O(n)$ algorithm for the 100 prisoners problem is precisely that a lower bound prevents us from doing so. The expected number of days for all prisoners to enter the interrogation room at least once is $O(n\log n)$, therefore no strategy, no matter how clever, may achieve a better expected runtime than $O(n\log n)$, [1]. A simple calculation confirms this lower bound. Let the random variable X_i be the number of days until the i-th unique prisoner with probability of selection $\frac{n-i+1}{n}$ is picked.

$$E[X] = \sum_{i=1}^{n} E[X_i] = \sum_{i=1}^{n} \frac{n}{n-i+1}$$
$$= n\sum_{i=1}^{n}\frac{1}{i} = O(n\log n)$$

Naturally, when the original problem has been solved, we wonder if the solution still applies for variants of the problem. Some of these solutions, like the leader and the $O(n(\log n)^2)$ solutions, depend on certain characteristics of the problem like the ability to tell time. What if we took away these abilities? Below, I present some harder instances of the 100 Prisoners puzzle and challenge you to find more efficient solutions for them.

Variations of the 100 Prisoners and a Light Bulb Puzzle

We assume for Problems 1 and 4 that the following are true: All prisoners are allowed to discuss their strategy on the first day. On the next day, they are each placed in an isolated cell with a window. The interrogation room contains a single light bulb that is initially switched off.

1. **Blue and Red Cells**: Each isolated cell is either painted completely blue or completely red. In addition to declaring that all prisoners have been interrogated, a confident prisoner must also correctly state the number of prisoners in red cells and the number of prisoners in blue cells [1].

Problem 1 is easily solvable using a strategy similar to the "leader" strategy if two light bulbs are in the interrogation room instead of one. However, with only one light bulb, is it possible to devise an $O(n^2)$ time algorithm?

2. **Light Bulb May Be Off**: We assume that the light bulb in the interrogation room may be turned on or off initially (i.e. before the first prisoner enters) [2].

If prisoners still have windows in their rooms, then the "leader" algorithm still provides an $O(n^2)$ solution to Problem 2 because the leader can just record all the times the light is on starting from the second day. The first non-leader prisoner to enter the interrogation room must be unique; therefore, on the first day, he can simply leave the light on if it is on or turn it on if it is off. All other prisoners behave as before. However, this problem becomes trickier if prisoners do not have windows in their individual cells because the prisoners have just lost their ability to keep track of time.

3. No Windows: We keep the condition presented in Problem 2. Now, prisoners may no longer keep track of how much time has passed because they are placed in isolated cells with no windows and no way to keep time [2].

This variation is harder because now the leader does not know how many days have passed and how many prisoners were interrogated before she enters the room. She could be the first prisoner to enter the room and the light bulb could have been initially on. In this case, her count of the number of interrogated prisoners would be off by 1. Can we still achieve an $O(n^2)$ algorithm by tweaking the "leader" protocol (the answer is yes but how)? The harder question is can we tweak the $O(n(\log n)^2)$ solution to apply to this problem?

4. **Couple of Prisoners**: Let us assume that all prisoners arrested were couples. Therefore, among the 100 prisoners, there are 50 distinct couples (no person may be a member of more than one couple). The warden then divides each couple. One member of the couple is placed in Group A and the other is placed in Group B. On each day, the warden chooses uniformly at random with replacement someone in Group A to interrogate in the morning. In the afternoon, on the same day, the warden chooses randomly someone from Group B to interrogate. Couples may not switch who they're partnered with. In addition to declaring that all 100 prisoners have been interrogated, a prisoner must also correctly claim that all couples have been interrogated (at least once) on the same day [4].

There exists a solution that assigns each couple to a particular day. The person from Group A may only turn the light on when they are called on their assigned day. Otherwise, they turn the light off. If the person from Group B is also called on their assigned day, they will leave the light on if it is on from the morning. If the person from Group B is called on any other day, they will turn the light off. A leader chosen from Group A counts the number of unique days she sees a light on when she enters the room. This indicates that both members of a couple were interrogated on their assigned day (the previous day). What is the expected runtime of this solution? Does this problem still have a solution if the prisoners are placed in isolated cells without windows?

Trading Light Bulbs for Time

For this last problem, I want to see how much more power we must give to the prisoners in order to bring the expected number of days in jail down to $O(n)$. The riddle I created below trivially must have an $O(n)$ solution (answer in the appendix). Giving prisoners more light bulbs enables them to tell each other more information in a shorter amount of time. But the more interesting question, now, is can the prisoners escape with less than 6 light bulbs?

5. **Prisoners and Vindictive Wardens:** The same 100 prisoners are ushered into prison by the same warden. They will be placed in isolated cells with windows starting tomorrow. Except now, the warden tells them that the interrogation room has 6 light bulbs in a row, and she will interrogate each prisoner at most twice. Prisoners are chosen uniformly at random from those that have not yet been interrogated twice. The prisoners were tremendously happy at this news because they are guaranteed freedom after at most 200 days. The warden cackles and tells them that there is a catch. This time, when a prisoner enters the interrogation room, he is asked, "Are you the last unique prisoner?" The last unique prisoner must declare, "Yes." Every other prisoner must declare, "No." Once a "Yes" is correctly declared, everyone is immediately freed. If someone declares incorrectly, everyone will be executed. How can they guarantee their freedom?

As a hint, the solution to this problem critically depends on the prisoners being able to tell time. If the isolated cells do not contain windows, what, then, is the minimum number of light bulbs needed in order to guarantee the prisoners' freedom?

Is the Solution Optimal?

I hope that you will take what I have written here to heart so that the next time you look at a puzzle, don't just find a right solution; find the optimal solution.

Appendix

Answer to "Prisoners and Vindictive Wardens": Despite having 6 light bulbs, the answer to this riddle is not as simple as encoding the number of prisoners who have been interrogated twice, which we could if we had 7 light bulbs. But we may use a similar scheme. Let "on" represent 1 and "off" represent 0, with the rightmost light bulb representing the smallest bit. On the i-th day, the prisoner who enters the interrogation room knows at least $\lfloor \frac{i}{2} \rfloor$ unique prisoners must have already been interrogated by the pigeonhole principle. Then, using this fact and the light bulbs, we may implement a counting system. Let Δ_i be the number represented in bits by the 6 light bulbs on the i-th day. Every prisoner knows how many times he has been interrogated. If the prisoner is entering the interrogation room the first time, he will check whether $\lfloor \frac{i}{2} \rfloor + \Delta_i$ is 99. If so, then he declares, "Yes." If not, he changes the light bulbs such that $\lfloor \frac{i+1}{2} \rfloor + \Delta_{i+1} = \lfloor \frac{i}{2} \rfloor + \Delta_i + 1$ and declares "No." If the prisoner is entering the room for the second time, he will change the light bulbs such that $\lfloor \frac{i+1}{2} \rfloor + \Delta_{i+1} = \lfloor \frac{i}{2} \rfloor + \Delta_i$ and declare "No." We may see this algorithm works for any n prisoners because $0 \leq \Delta_i \leq \lfloor \frac{n}{2} \rfloor + 1$. Never is $\Delta_i > \lfloor \frac{n}{2} \rfloor + 1$ because that means $\lfloor \frac{n}{2} \rfloor + \Delta_i > n$, a contradiction. Furthermore, never is $\Delta_i < 0$ because we would contradict the pigeonhole principle. For any n prisoners, this scheme would work given $\lceil \log_2 \left(\lfloor \frac{n}{2} \rfloor + 1 \right) \rceil$ light bulbs. May we achieve a better scheme using fewer light bulbs?

References

1. P.O. Dehaye, D. Ford, and H. Segerman, One hundred prisoners and a lightbulb, Mathematical Intelligencer (2003).

2. H. van Ditmarsch, J. van Eijck, and W. Wu, Verifying One Hundred Prisoners and a Lightbulb, Journal of Applied Non-Classical Logics (2007).

3. W. Wu, 100 prisoners and a lightbulb, https://www.ocf.berkeley.edu/~wwu/papers/100prisonersLightBulb.pdf (2002).

4. http://wiki.scn.sap.com/wiki/pages/viewpage.action?pageId=360972998.