100 PRISONERS AND A LIGHT BULB

WILLIAM WU

Abstract. We present a variety of different protocols for solving the "100 Prisoners and a Light Bulb" riddle, including explicit computations of average runtimes.

This is an updated and corrected version of the original unpublished document released on December 5, 2002 at `http://www.ocf.berkeley.edu/~wwu/`. The protocols presented here can be attributed to many members of [wu::forums], an online forum of mathematical riddles at `wuriddles.com`. Our quest to optimize the 100 prisoners problem is never ending, and there are many clever algorithms that have yet to be incorporated in this article.

Contents

1. The Riddle

One hundred prisoners have been newly ushered into prison. The warden tells them that starting tomorrow, each of them will be placed in an isolated cell, *unable to communicate amongst each other. 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 prisoner will be able to observe the current state of the light bulb. If he wishes, he can 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 protocol that will guarantee their freedom?

It may seem surprising at first that such a protocol could exist. We will present an assortment of such protocols that guarantee freedom, and analyze and compare the average number of days that each protocol requires. To compare the protocols fairly, we will express their average runtimes in terms of N, where n is the number of prisoners in general.

Before we begin, some preliminary remarks. Unless otherwise stated, we will make the following assumptions throughout this paper:

(1) Prisoners can count how may days have elapsed.
(2) The initial bulb state is OFF.

The second assumption is a trivial consequence of the former, since we can have the prisoner who enters on the first day turn the bulb OFF. However, dropping the first assumption would render many of the more advanced protocols infeasible.

2. Contributions and Related Work

The objective of this work is to systematically collect and analyze all known algorithms for solving the 100 Prisoners and Light Bulb riddle. The 2004 Mathematical Intelligencer article by Dehaye, Ford and Segerman [1] also discusses solution protocols for 100 Prisoners and Light Bulb. Their article discusses many interesting variants of the riddle (13 out of 18 pages). Our main contribution, in contrast, is to provide detailed probabilistic analyses for all known solutions to the original riddle. We also correct some errors in the description of the binary tokens scheme presented in [1], and mention a few protocols not discussed there.

The protocols presented here could be credited to the combined efforts of many dedicated members of `[wu::forums]`, an online community forum of mathematical riddles at `wuriddles.com`.

3. Origins of the Riddle

Unfortunately the origins of this riddle are unclear, so the author can only discuss his personal experiences. He first heard about it in 2001, through members of the University of California Berkeley Chapter of Eta Kappa Nu (HKN), an Electrical Engineering Honor Society. Puzzles are often circulated in HKN because new students seeking to be inducted into the society are required to complete a series of challenges, and often these challenges come in the form of logical or mathematical puzzles. Later, in 2002, the author found the problem also listed on one of the challenges webpage of IBM Research [2], wherein it was mentioned that "this puzzle has been making the rounds of Hungarian mathematicians' parties".

Here is a whimsical piece of history. When the author posted the riddle in 2002 on his website `wuriddles.com`, he add some fictional story line of his own. Namely, he wrote that if the assertion that all prisoners have been in the room is false, the prisoners would be "shot for their stupidity", whereas if the assertion is true, the prisoners are "set free and inducted into MENSA, since the world could always use more smart people." And due to the "slashdot effect", this wording of the puzzle involving MENSA can now be found in many places. In retrospect, the qualifier clause for MENSA is not very logical, since nothing could be more dangerous than criminally-minded smart people. In any case, if you see a wording of the 100 prisoners problem that involves MENSA, you now know who added that nonsense.

4. I'm Feeling Lucky Protocol

4.1. **Protocol.** The days are split into n-day blocks. During each n-day block, each prisoner operates according to the following instructions upon entering the interrogation room:

- If it is day 1 for the current block:
  - If the bulb is OFF, turn the bulb ON.
  - If the bulb is already ON, and the first n-day block has already elapsed, announce that all prisoners have visited.
- On any other day of the current block:
  - If it is your first time visiting the room during the current block, do nothing.
  - If it is your second time visiting the room during the current block, turn the light OFF.
  - If it is your third or more time visiting the room during the current block, do nothing.

The general idea is that eventually, with probability 1, we will be lucky enough to have a block of n-days during which no prisoner enters the room twice, or in other words, during which every prisoner will enter the room exactly once. Then the bulb which was turned ON on day 1 will still be ON after n-days, since bulbs are only turned OFF upon a second return visit. Thus, if the bulb remains ON on the first day of a new block, we know that every prisoner must have visited the interrogation room during the block that had just elapsed.

4.2. **Expected Runtime.** We now compute the expected runtime of this protocol. Let X be the number of days the protocol requires. Let B be the number of n-day blocks required till the protocol succeeds. Then B is a geometric random variable with parameter

$$\frac{n}{n} \cdot \frac{n-1}{n} \cdot \frac{n-2}{n} \cdot \frac{n-3}{n} \cdots \frac{1}{n} = \frac{n!}{n^n}.$$

Since the expectation of a geometric random variable is the reciprocal of its parameter, and $X = nB$, the expected number of days required is

$$\mathbf{E}\left[X^{(\text{i'm feeling lucky})}\right] = n\mathbf{E}\left[B\right] = n\frac{n^n}{n!} = \frac{n^{n+1}}{n!}. \tag{1}$$

Using Stirling's approximation $n! \sim \sqrt{2\pi n}\left(\frac{n}{e}\right)^n$, and big-O notation,

$$\mathbf{E}\left[X^{(\text{i'm feeling lucky})}\right] \sim \frac{1}{\sqrt{2\pi}} n^{1/2} e^n = O(n^{1/2}e^n). \tag{2}$$

When $n = 100$, $\mathbf{E}\left[X\right]$ equals 1.072×10^{44} days. We can expect the prisoners to be long dead by then, and thus, it would behoove us to design a faster protocol.

5. Single Counter Protocol

5.1. **Protocol.** One of the possible sources of difficulty in solving this riddle is the natural idea that every prisoner should should follow the same instructions. I will call such a protocol *symmetric*; the previous protocol is an example of such. Realizing that such a constraint does not actually exist, we can design a simple asymmetric scheme that performs much better, allowing the prisoners to escape within a reasonable amount of time on average.

Letting prisoners have different roles, we assign one prisoner to be "the counter". He will maintain an integer variable in his head that is initialized to 1. Call this variable T. Upon entering the room, prisoners adhere to the following instructions:

- If you are not the counter:
  - If the bulb is OFF, and you have never turned the bulb ON before, turn it ON.
  - If the bulb is ON, do nothing.
- If you are the counter:
  - If the bulb is OFF, do nothing.
  - If the bulb is ON, turn it OFF, and set T=T+1.
  - If T = n, announce that all prisoners have visited.

The idea behind this protocol is that every prisoner besides the counter will turn ON the bulb exactly once, whenever he can. When the bulb is ON, no one can turn it OFF except for the counter. Eventually the counter will enter the room, turn this bulb OFF, and increment the count T. In this way, each prisoner indicates his presence in the room to the counter by leaving an ON bulb which is eventually recorded by the counter.

5.2. **Expected Runtime.** To analyze the runtime, we can split the process into epochs. Let X_i denote the number of days between the first day on which $T = i$, and the first day on which $T = i + 1$. Between these two days, two events must occur:

1. An unrecorded prisoner must be chosen, causing the bulb to be turned ON. Let Y_i denote the number of days between from when $T = i$ until this event occurs.
2. The counter must then enter the room to record this ON bulb. Let Z_i denote the number of days from when the bulb is turned ON until this occurs.

Then

$$X_i = Y_i + Z_i.$$

Letting X be a random variable corresponding to the number of days the protocol requires in total, we have

$$X = \sum_{i=1}^{n-1} X_i = \sum_{i=1}^{n-1} (Y_i + Z_i).$$

Y_i is a geometric random variable with parameter $\frac{n-i}{n}$, and Z_i is a geometric random variable with parameter $\frac{1}{n}$. Hence, by linearity of expectation, the expected runtime is

$$\begin{aligned} \mathbf{E}\left[X^{(\text{one counter})}\right] &= \sum_{i=1}^{n-1} (\mathbf{E}[Y_i] + \mathbf{E}[Z_i]) \\ &= \sum_{i=1}^{n-1} \left(\frac{n}{n-i} + n\right) \\ &= (n-1)n + n\sum_{i=1}^{n-1} \frac{1}{i} \\ &= n^2 - n + nH_{n-1}. \end{aligned} \tag{3}$$

In big O, since $H_n \sim \ln n$,

$$\mathbf{E}\left[X^{(\text{one counter})}\right] = O(n^2). \tag{4}$$

When $n = 100$, $\mathbf{E}[X]$ equals 10417.74 days, or 28.54 years, which is still within the span of a young prisoner's lifetime.

The variance of this protocol may also be easily computed. Since the variance of a sum of independent random variables is the sum of the variances, and the variance of a geometric random variable with parameter p is $\frac{1-p}{p^2}$,

$$\begin{aligned} \mathbf{var}(X^{(\text{one counter})}) &= \sum_{i=1}^{n-1} (\mathbf{var}(Y_i) + \mathbf{var}(Z_i)) \\ &= \sum_{i=1}^{n-1} \left(\frac{1-\left(\frac{n-i}{n}\right)}{\left(\frac{n-i}{n}\right)^2} + \frac{1-\frac{1}{n}}{\left(\frac{1}{n}\right)^2}\right) \\ &= \sum_{i=1}^{n-1} \left(\frac{ni}{(n-i)^2} + (n-1)n\right) \\ &= n(n-1)^2 + n\sum_{i=1}^{n} \frac{i}{(n-i)^2} \\ &= n(n-1)^2 + n\sum_{j=1}^{n} \left(\frac{n}{j^2} - \frac{1}{j}\right) \\ &= n(n-1)^2 + n(nH_{n-1,2} - H_{n-1}) \end{aligned} \tag{5}$$

where $H_{n,2} := \sum_{i=1}^{n} \frac{1}{i^2}$. Asymptotically,

$$\mathbf{var}(X^{(\text{one counter})}) = O(n^3). \tag{6}$$

The one counter protocol is the "standard solution" to the puzzle. In the sequel, we will describe some less well-known solutions that perform even better.

6. One Counter Protocol, with Non-Counters Feeling Lucky

6.1. **Protocol.** Under the one counter protocol, the prisoners escape if and only if the bulb, which is initially OFF, alternates its state from OFF to ON exactly $n-1$ times. Non-counters can also count these state transitions as they witness them. So, a marginal improvement in the algorithm can be made by realizing that if any very lucky non-counter witnesses all $n-1$ such transitions before the counter does, then the non-counter is equally qualified to declare victory and preempt the counter in the very last epoch of the algorithm.

Since the standard one counter protocol already requires a runtime of $O(n^2)$, and this new policy for non-counters can only save at most n days (since it only affects the last epoch), the improvement does not affect the asymptotics, so we will not mention it again. Furthermore, the probability of a non-counter declaring victory under this scheme approaches zero extremely quickly as n grows. The author's explicit computations for this probability are a bit long, so they are left to Appendix B.

7. One Counter Protocol, with Dynamic Counter Assignment

7.1. **Protocol.** The One Counter Protocol can be slightly improved by assigning the role of counter dynamically, rather than a priori. We use the following policy: the counter is the first person to enter the room twice in the first n days.

- Stage I: Days 1 through n:
  - Days 1 through $n-1$: The first person to enter the room twice will turn the bulb ON, and assign himself to be the counter.
  - Day n: If the light is still OFF, declare victory. Otherwise, turn off the light.
- Stage II: (all remaining days)
  Follow the normal One Counter Protocol, but with the following modifications:
  - The counter only counts up to $n-k+1$, where k is the index of the day that the counter entered the interrogation room twice.
  - Prisoners who saw an ON bulb in Stage I do nothing.

To illustrate the idea behind this protocol, suppose we have 100 prisoners, and the first person to enter the interrogation room twice enters on day 20. This prisoner becomes the counter, and he can deduce that in the previous 19 days, there have been exactly 19 distinct visitors, including himself. Thus, when Stage II ensues, he would only need to tally (n-1) - (k-2) = n-k+1 = 99-18 = 81 prisoners. Lastly, if we are so lucky that no counter is assigned on the 100th day, then every visitor in the first 100 days must have been distinct (as in the I'm Feeling Lucky scheme), so we declare victory.

7.2. **Expected Runtime.** We now compute the expected runtime of this modified protocol. Let the random variable X represent the total number of days till victory is declared, and let the random variable K represent the day on which a prisoner first re-enters the room. By the total probability theorem,

$$\mathbf{E}[X] = \sum_{k=2}^{n+1} \mathbf{P}[K = k]\, \mathbf{E}[X|K = k]. \tag{7}$$

We now compute each of these terms. First we will compute $\mathbf{E}[X|K = k]$. By the pigeonhole principle, the largest possible value of K is $n+1$. In the special case that $K = n+1$, we should have $X = n$, since we will declare victory on day n. Excluding this case for now, suppose $K \leq n$. Then, in Stage II, the counter must count up to $n - K + 1$. To see this, observe that in the first K days, there are $K - 1$ distinct prisoners including the counter himself, and thus $K - 2$ prisoners which the counter will not have to count in Stage II. Since the counter normally has to account for $n - 1$ other prisoners, this leaves $n - 1 - (K - 2) = n - K + 1$ remaining that are unaccounted for. $\mathbf{E}[X|K = k]$ can then be broken into epochs, just as we did in analyzing the One Counter Protocol. We first add n days, since that is the fixed length of Stage I. Then, starting in Stage II, the number of days till the first unaccounted prisoner enters the room and turns ON the bulb is a geometric random variable with parameter $\frac{n-k+1}{n}$. Afterwards, the number of days till the counter records this bulb is a geometric random variable with parameter $\frac{1}{n}$. The number of days till the second unaccounted prisoner enters is then a geometric random variable with parameter $\frac{n-k}{n}$, and so forth. Thus,

$$X|\{K = k\} = n + \sum_{i=1}^{n-k+1} (Y_i^{(k)} + Z_i^{(k)})$$

where $Y_i^{(k)} \sim \mathbf{geom}\left(\frac{n-k+2-i}{n}\right)$ and $Z_i^{(k)} \sim \mathbf{geom}\left(\frac{1}{n}\right)$ for $i \in \{1, 2, \ldots, n-k+1\}$. Taking expectation,

$$\begin{aligned}
\mathbf{E}[X|K = k] &= n + \sum_{i=1}^{n-k+1} \left(\mathbf{E}\left[Y_i^{(k)}\right] + \mathbf{E}\left[Z_i^{(k)}\right]\right) \\
&= n + \sum_{i=1}^{n-k+1} \left(\frac{n}{n-k+2-i} + n\right) \\
&= n + n(n-k+1) + n \sum_{i=1}^{n-k+1} \frac{1}{n-k+2-i} \\
&= n + n(n-k+1) + n \sum_{i=1}^{n-k+1} \frac{1}{i} \\
&= n(2 + n - k + H_{n-k+1}).
\end{aligned} \tag{8}$$

Observe that this formula also works in the fringe case that $K = n+1$, in which case

$$\mathbf{E}[X|K = n+1] = n(2 + n - n - 1 + H_0) = n.$$

Secondly, we must compute $\mathbf{P}[K = k]$. Comparing each draw of a random prisoner to a ball thrown into one of n bins, $\mathbf{P}[K = k]$ is the probability that the first bin collision occurs on the k^{th} throw. Stepping through the process, the probability that the first throw has no collision is 1. Afterwards, the probability that the second throw has no collision is $\frac{n-1}{n}$, and the probability that the third throw does not collide with either of the first two throws is $\frac{n-2}{n}$. These misses continue until the $(k-1)^{th}$ throw. The probability of the k^{th} throw then being the first collision is then $\frac{k-1}{n}$. Hence,

$$\mathbf{P}[K = k] = 1 \cdot \frac{n-1}{n} \cdot \frac{n-2}{n} \cdots \frac{n-(k-2)}{n} \cdot \frac{k-1}{n} = \frac{n^{\underline{k-1}}}{n^k}(k-1). \tag{9}$$

Note that since a probability mass function sums to 1, we apparently have $1 = \sum_{k=2}^{n+1} \frac{n^{\underline{k-1}}}{n^k}(k-1)$. We can rewrite this as the following identity, which we will use in the future:

$$n = \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k. \tag{10}$$

Plugging Equations 8 and 9 into Equation 7, we have

$$\begin{aligned} \mathbf{E}\left[X^{(\text{dynamic one counter})}\right] &= \sum_{k=2}^{n+1} \mathbf{P}[K = k]\, \mathbf{E}[X | K = k] \\ &= \sum_{k=2}^{n+1} \frac{n^{\underline{k-1}}}{n^k}(k-1) \cdot n(2 + n - k + H_{n-k+1}) \\ &= \sum_{k=2}^{n+1} \frac{n^{\underline{k-1}}}{n^{k-1}}(k-1) \cdot (2 + n - k + H_{n-k+1}) \\ &= \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(1 + n - k + H_{n-k}) \\ &= \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(1 + n) + \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(H_{n-k} - k) \\ &= n^2 + n + \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(H_{n-k} - k) \end{aligned} \tag{11}$$

where in the last equality we have used Equation 10. The asymptotics of this expression can be crudely upper bounded to be no worse than those of the original one-counter scheme:

$$\begin{aligned} \mathbf{E}[X] &\leq n^2 + n + \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(H_n - 0) \\ &= n^2 + n + (\ln n) \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k \\ &= n^2 + n + n \ln n. \end{aligned} \tag{12}$$

Using some delicate bounding, one can also show that the big O runtime of this modified protocol is actually strictly less than n^2, and thus asymptotically superior. The author has been unable to find a short proof of this, so for the sake of brevity, the proof is omitted here.

However, using Equations 3 and 11, it does not take more than a page to analytically show that dynamic counter assignment does constitute an improvement in average runtime over the One Counter scheme:

$$\begin{aligned}
\mathbf{E}\left[X^{(\text{one counter})}\right] - \mathbf{E}\left[X^{(\text{dynamic one counter})}\right] &= n^2 - n + nH_{n-1} - n^2 - n - \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(H_{n-k} - k) \\
&= nH_{n-1} - 2n + \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(k - H_{n-k}) \\
&\geq nH_{n-1} - 2n + \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k(k - H_{n-1}) \\
&= nH_{n-1} - 2n + \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k^2 - H_{n-1} \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k \\
&= nH_{n-1} - 2n + \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k^2 - nH_{n-1} \\
&= \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k^2 - 2n \\
&= \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k^2 - 2 \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} k \\
&= \sum_{k=1}^{n} \frac{n^{\underline{k}}}{n^k} (k^2 - 2k) \\
&= -1 + \sum_{k=2}^{n} \frac{n^{\underline{k}}}{n^k} (k^2 - 2k) \\
&= -1 + \sum_{k=3}^{n} \frac{n^{\underline{k}}}{n^k} (k^2 - 2k).
\end{aligned}$$

The terms in the summation are all positive. Assuming $n \geq 4$, we can lower bound the first term in the summation by

$$\begin{aligned}
\frac{n^{\underline{k}}}{n^k}(k^2 - 2k)\Big|_{k=3} &= 6\frac{n^{\underline{3}}}{n^3} \\
&= 6\frac{n(n-1)(n-2)}{n^3} \\
&= 6\left(1 - \frac{1}{n}\right)\left(1 - \frac{2}{n}\right) \\
&\geq 6\left(1 - \frac{1}{4}\right)\left(1 - \frac{2}{4}\right) \\
&= 9/8.
\end{aligned}$$

Thus, for all $n \geq 4$,

$$\mathbf{E}\left[X^{(\text{one counter})}\right] - \mathbf{E}\left[X^{(\text{dynamic one counter})}\right] \geq 1/8 > 0.$$

8. Two Stage Counting Protocol

8.1. **Protocol.** Observe that in the Single Counter Protocol, we will have long stretches of time where the bulb is ON and we are waiting for the counter to enter the room. This suggests that it maybe useful to have multiple alternative counters who are also authorized to record the ON bulb and turn it OFF. Furthermore, it would be nice if we could count faster to n. That is, rather than counting 1-by-1 to n, what if we counted in jumps of 10 instead?

The Two-Stage Counting Protocol improves on the Single Counter Protocol in both of the aforementioned aspects. Firstly, it divvies up the task of counting the prisoners amongst a group of assistant counters. Secondly, the head counter counts up to n more quickly by collecting the aggregated counts of the assistant counters.

To begin the protocol's description, there are three different possible roles for a prisoner: head counter, assistant counter, and "drone". There is exactly one head counter, and there is some number of assistant counters $a \ll n$, while the vast majority prisoners are still drones – regular prisoner with no counting tasks. The head counter and all assistant counters all have an integer variable in their heads, initialized to one.

The protocol has two stages, Stage I and Stage II. Each stage lasts for a certain number of preset days, which we will call s_1 and s_2, respectively. In Stage I, each assistant counter is responsible for counting a quota of q drones. In Stage II, the head counter will be responsible for counting up the assistant counters who have reached their quota. In this way, the head counter counts toward n in jumps of size q. If the head counter does not succeed by the end of Stage II, then we repeat Stage I and Stage II again, still maintaining all the mental counts from before. In other words, we repeatedly alternate between Stages I and II until victory is declared. Notice that there were five parameters required by the protocol: n, a, q, s_1, and s_2.

Getting down to brass tacks, upon entering the interrogation room, each prisoner adheres to the following instructions:

- During the first $s_1 - 1$ days of Stage I:
  - If you are a drone:
    * If the bulb is OFF, and you have not turned it ON before, turn it ON.
    * If the bulb is ON, do nothing.
  - If you are an assistant counter:
    * If the bulb is OFF, do nothing.
    * If the bulb is ON, and you have not reached your quota yet, turn it OFF, and increment your count.
  - If you are the head counter: always do nothing.
- On the last day of Stage I:
  - If you are a drone:
    * If the bulb is OFF, do nothing.
    * If the bulb is ON, turn it OFF, and plan to turn the bulb ON one extra time in future invocations of Stage I.
  - If you are an assistant counter:
    * If the bulb is OFF, do nothing.
    * If the bulb is ON, turn it OFF, and increment your count.
  - If you are the head counter:
    * If the bulb is OFF, do nothing.
    * If the bulb is ON, turn it OFF, and plan to turn the bulb ON one extra time in future invocations of Stage I.
- During the first $s_2 - 1$ days of Stage II:
  - If you are a drone: always do nothing.
  - If you are an assistant counter:
    * If the bulb is OFF, and you have reached your quota q and have not turned it ON before, turn it ON.
    * If the bulb is ON, do nothing.
  - If you are the head counter:
    * If the bulb is OFF, do nothing.
    * If the bulb is ON, turn it OFF, and increment your count by q.
    * If your count equals n, declare victory.
- On the last day of Stage II:
  - If you are a drone:
    * If light is ON, turn it OFF. Then in future invocations of Stage I, turn the light ON q times.
    * If light is OFF, do nothing.
  - If you are an assistant counter:
    * If light is ON, turn it OFF. Then in future invocations of Stage II, turn the light ON once. (This task is in addition to the default task that each assistant counter has. So this assistant counter may have to turn on the light more than once in future Stage IIs.)
    * If light is OFF, do nothing.
  - If you are the head counter:

  - $*$ If the light is OFF, do nothing.
  - $*$ If the light is ON, add q to count. If the count equals n, declare victory.

The algorithm is a little more complex than one might expect due to the care that must be taken on the last days of Stages, to assure that "no light is left behind." I am very indebted to Hans van Ditmarsch for revealing many of these special cases to me.

8.2. **Expected Runtime.** The average runtime of this algorithm is difficult to compute, and remains open for now. However, simulations with certain parameters for the case of $n = 100$ yield runtimes between 3500 and 4000 days, or 9.5 to 11 years.

9. Binary Tokens Protocol

9.1. **Protocol.** The basic idea behind the two stage counting protocol was that to speed things up, sometimes we should count in clumps rather than one-by-one. In the first stage, assistant counters counted one-by-one, and the second stage, the master counter counted the clumps collected by the assistant counters.

This same protocol can be thought of in terms of exchanging "tokens" with variable point values. To make the analogy clear, imagine that all prisoners not assigned any counting roles start with a token worth one point. During Stage 1, these prisoners try to deposit their one-point tokens into the central room by turning on the bulb when they can, and assistant counters collect the tokens. Suppose assistant counters are ordered to count up to 10. Then in Stage 2, assistant counters exchange their collected tokens with 10-point tokens, and try to deposit these 10-point tokens into the room by turning on the bulb when they can. The master counter collects these bigger tokens. Thus, a lighted bulb represents a different number of points depending on what stage we are in, and the prisoners can escape more quickly by counting in terms of gradually higher denomination tokens.

The "binary tokens scheme" is a generalization of these ideas. The value of a lighted bulb is doubled from stage to stage, and all prisoners now have the same role, allowed both to deposit points and collect points. Proceeding formally, let n be the total number of prisoners, and suppose n is a power of 2. Let P_k be the number of points a lighted bulb is worth on day k. We will define it later, but for now, know that every P_k is a nonnegative power of 2. All prisoners use the following instructions:

- Keep an integer in your head; call it T. Initialize it to $T = 1$.
- Let T_m denote the m^{th} bit of T expressed in binary.
- Upon entering the room on day k, where $P_k = 2^m$ for some m, go through four steps:
  1. If the bulb is ON, set $T := T + P_{k-1}$, and turn it OFF.
  2. If $T \geq n$, declare victory.
  3. If $T_m = 1$, turn the bulb ON, and set $T := T - P_k$.
  4. Else, if $T_m = 0$, leave the bulb OFF and do nothing.

Notice that Step 1 amounts to taking a token worth P_{k-1} points left over from the previous day, and Step 2 amounts to depositing a token worth P_k points. In short, all prisoners will collect and deposit tokens whenever they may legally do so, where the value of tokens are universally dictated by a prespecified sequence P_k that is *only* a function of what day it is. Whenever someone accumulates 100 points worth of tokens, the game is over.

It remains to specify what (P_k) should be. The sequence should start with a block of consecutive ones, since everyone starts with only one point. If this block is long enough, there will be many prisoners who have collect more than one point, and perhaps a subsequent block of twos would be effective. For reasons that will become apparent in the coupon-collection analysis presented in the following section, we choose the nondecreasing sequence

$$(P_k : k \in [1:T]) = (\underbrace{1,1,\ldots,1}_{n\ln n + n\ln\ln n}, \quad \underbrace{2,2,\ldots,2}_{n\ln n + n\ln\ln n}, \quad \underbrace{4,4,\ldots,4}_{n\ln n + n\ln\ln n}, \ldots, \underbrace{\frac{n}{2},\frac{n}{2},\ldots,\frac{n}{2}}_{n\ln n + n\ln\ln n})$$

where $T := \log_2 n(n\ln n + n\ln\ln n)$, the length of the finite sequence on the right-hand side. There are $\log_2 n$ stages, each lasting $n\ln n + n\ln\ln n$ days (rounded). In the k^{th} stage, the bulb is worth 2^k, where k indexes from 0 to $(\log_2 n) - 1$.

Lastly, if victory has not been declared after T days, the prisoners will maintain the integers in their heads, and (P_k) restarts. That is, the full sequence (P_k) is T-periodic:

$$P_k := 2^m \qquad \text{where} \qquad m := \left\lfloor \frac{k \pmod T}{n\ln n + n\ln\ln n} \right\rfloor.$$

9.2. **Expected Runtime.** The goal of this section is to prove that the average runtime of the binary tokens protocol is

$$O(n(\ln n)^2).$$

To outline our approach, we will first show that the binary tokens protocol can be reduced to a succession of coupon collector problems. (Recall the coupon collector problem: suppose there are n different possible coupons, and each day we receive one of them uniformly at random. The objective is then to collect all n coupons.) After having done so, we will modify the proof of the following well-known result to suit our needs (see Appendix A):

Lemma 1. *In a coupon collection problem with n coupons, after $n\ln n + cn$ draws, the probability of not having seeing all the coupons is less than $\frac{1}{e^c}$.*

Working through a simple example will illustrate why our problem is related to coupon collection. Suppose we have $n = 4$ prisoners labeled A, B, C, and D. Stage 0, in which the bulb is always worth 1 point, then lasts for $\lceil n\ln n + n\ln\ln n\rceil$ days. In the beginning, every prisoner starts with one point, and the bulb is OFF. We can represent this initial state by the table

Day 0:

OFF	2^1	2^0
A	0	1
B	0	1
C	0	1
D	0	1

where the bulb's status is indicated in the upper left, and the integers being mentally maintained by each of the prisoners is listed in binary in the lower right. Let us play out the following sequence of visitations in Stage 0: $A, B, C, B, A, \ldots, D$.

On Day 1, A is chosen. Following the protocol, A will turn the bulb ON and decrement his number. The new state becomes:

End of Day 1:

ON	2^1	2^0
A	0	0
B	0	1
C	0	1
D	0	1

On Day 2, B is chosen. He sees the ON bulb, turns it off, and increments his count. He then checks if the zeroth bit of his newly incremented count is a 1, but it is not, so he does not activate the bulb. The new state is:

End of Day 2:

OFF	2^1	2^0
A	0	0
B	1	0
C	0	1
D	0	1

On Day 3, C is chosen. This leads to:

End of Day 3:

ON	2^1	2^0
A	0	0
B	1	0
C	0	0
D	0	1

On Day 4, suppose that B is chosen again. B sees the bulb, still worth 1 point, and turns it OFF. He then increments his count to $2+1=3$, which is 11_2 in binary. Then he sees that the zeroth bit of his count so far is a 1, so he decrements his count back to 2, and turns the bulb ON again. So within Day 4, we have

Start of Day 4:

OFF	2^1	2^0
A	0	0
B	1	1
C	0	0
D	0	1

$\longrightarrow$ End of Day 4:

ON	2^1	2^0
A	0	0
B	1	0
C	0	0
D	0	1

which is the same state as the previous day. In short, choosing B again has no effect on the system.

Now suppose that on Day 5, A (or equivalently, C) is chosen. The consequent behavior will again be identical to that of B on Day 4. Any prisoner with a zeroed count will add and then immediately subtract out whatever the bulb is worth on that day to his count, resulting in no net state change. Thus, any prisoner whose count reaches zero can be thought of as being *inactive* for the rest of this stage.

Hence, we see that in the remaining days of Stage 0, no net state change will occur unless D, the only person unchosen so far, is chosen, which would lead to the last state in Stage 0:

$$\begin{array}{c|cc} ON & 2^1 & 2^0 \\ \hline A & 0 & 0 \\ B & 1 & 0 \\ C & 0 & 0 \\ D & 0 & 1 \end{array} \quad \longrightarrow \quad \begin{array}{c|cc} OFF & 2^1 & 2^0 \\ \hline A & 0 & 0 \\ B & 1 & 0 \\ C & 0 & 0 \\ D & 1 & 0 \end{array}$$

Notice all ones have been paired into groups of two. Stage 1 then proceeds much like Stage 0 did, except that now we increment/decrement starting with the left column of bits, and the number of active prisoners has been halved from four to two. It is easy to see where this binary pattern is going; at the start of Stage k, we should have combined all the 2^{k-1} tokens into 2^k tokens, and there should be only $n/2^k$ active prisoners left.

When does the protocol fail? Notice that if D is never chosen in Stage 0, he will never have another chance to deposit his 1-point token into the room since the value of the bulb only goes up in future stages. Thus, the only way the protocol could succeed in this cycle (going through all stages once) is if D is the victory-declaring prisoner which collects all n points in the end. In general though, if there are ever even just two prisoners who are not chosen in a stage, this entire cycle is destined to fail. So, we can draw the following conclusion:

> Up to a negligible fencepost error, the binary tokens protocol succeeds if and only if in each stage, every active prisoner is chosen at least once, where the number of active prisoners in Stage k is $n/2^k$.

Thus each stage reduces to a coupon collection problem. In the k^{th} stage, we collect $n/2^k$ coupons (tokens), and we have $n\ln n + n\ln\ln n$ days to do it. Mimicking the proof of Lemma 1, if $\mathbf{P}\left[F_j^{(k)}\right]$ is the probability of failing to collect the j^{th} coupon at the k^{th} stage, where $j \in \{1, \ldots, n/2^k\}$, then

$$\begin{aligned} \mathbf{P}\left[F_j^{(k)}\right] &= \left(1 - \frac{1}{n/2^k}\right)^{n\ln n + n\ln\ln n} \\ &= \left(e^{-2^k}\right)^{\ln n + \ln\ln n} \qquad \text{as } n \to \infty \\ &= \left(e^{\ln n + \ln\ln n}\right)^{-2^k} \\ &= (n\ln n)^{-2^k} \\ &\leq \frac{1}{n\ln n}. \end{aligned}$$

Then, by invoking the union bound, $\mathbf{P}\left[F^{(k)}\right]$, the probability of the k^{th} stage failing, is

$$\mathbf{P}\left[F^{(k)}\right] = \mathbf{P}\left[\bigcup_{j=1}^{n/2^k} F_j^{(k)}\right] \leq \sum_{j=1}^{n/2^k} \mathbf{P}\left[F_j^{(k)}\right] \leq \frac{1}{2^k}\frac{1}{\ln n}.$$

Let F denote the event that one cycle of the protocol fails. Since the protocol fails if and only if at least one of the $\log_2 n$ stages fails, we can again invoke the union bound:

$$\mathbf{P}[F] = \mathbf{P}\left[\bigcup_{k=0}^{(\log_2 n)-1} F^{(k)}\right] \leq \sum_{k=0}^{(\log_2 n)-1} \mathbf{P}\left[F^{(k)}\right] \leq \sum_{k=0}^{(\log_2 n)-1} \frac{1}{2^k}\frac{1}{\ln n} \leq \frac{2}{\ln n}.$$

Let S denote the event that the first cycle of the protocol succeeds. Then

$$\mathbf{P}[S] = 1 - \mathbf{P}[F] \geq 1 - \frac{2}{\ln n}.$$

If the first cycle fails, then we can upper bound the probability that the second pass fails by the probability that the first cycle fails. This is true because the likelihood of successfully collecting all coupons at a given stage increases if some of these coupons were already collected in a previous cycle, thereby allowing for more opportunities for the uncollected coupons to be chosen. Thus, we can upper bound the expected number of cycles for the protocol by

$$\frac{1}{\mathbf{P}[S]} \leq \frac{1}{1-\frac{2}{\ln n}} \longrightarrow 1 \qquad \text{as } n \to \infty.$$

Hence, since each cycle consists of $\log_2 n$ stages, each of length $n \ln n + n \ln\ln n$, the total expected number of days till the prisoners get out is upper bounded by

$$\left(\frac{1}{1-\frac{2}{\ln n}}\right)(\log_2 n)(n \ln n + n \ln\ln n) \longrightarrow O(n(\ln n)^2).$$

10. Conclusion

We have presented five different protocols for solving the 100 Prisoners and a Light Bulb Riddle. The performance for each is summarized in Table 1, where n is the total number of prisoners.

Table 1. Summary of Protocols

protocol	closed-form average runtime	asymptotic average runtime
i'm feeling lucky	$\frac{n^{n+1}}{n!}$	$O(n^{1/2}e^n)$
one counter	$n^2 - n + nH_{n-1}$	$O(n^2)$
one counter dynamic	$n^2 + n + \sum_{k=1}^{n} \frac{n^k_-}{n^k} k(H_{n-k} - k)$	$\leq O(n^{2-\epsilon})$
two stages	n/a	n/a
binary tokens	n/a	$\leq O(n(\ln n)^2)$

In the future, we hope to present some lesser-known solution protocols developed in `[wu::forums]` [3] which, to our knowledge, have not been presented in any literature to date. We also hope to address how these protocols fare in variants of the 100 Prisoners and a Light Bulb riddle.

Appendix A. The Coupon Collector Problem

Suppose that there are n different possible coupons in the world, and each day we receive one of them uniformly at random in mail. We are then naturally interested in the following questions:

(1) What is the average number of days till we collect all n distinct coupons?
(2) Can we give a bound on the probability of collecting all coupons after m days?

To answer these questions, we first decompose the process into epochs. Let X be a random variable representing the total number of days until we see all N coupons. Let X_i be the number of days between first having seen $i-1$ distinct coupons and first having i coupons. Then X_i is a geometric random variable with parameter $\frac{n-i+1}{n}$, and $X = \sum_{i=1}^{n} X_i$. The expectation of X is then

$$\mathbf{E}[X] = \sum_{i=1}^{n} \mathbf{E}[X_i] = \sum_{i=1}^{n} \frac{n}{n-i+1} = n \sum_{i=1}^{n} \frac{1}{i} = nH_n \sim n \ln n.$$

One may notice that this analysis is identical to that of the one-counter protocol, which is a coupon collection problem as well.

For the second question, we will use the union bound. Suppose A is the event that not all coupons have been seen after m days have elapsed. Let A_i be the event that the i^{th} coupon has not been seen after m days. Then

$$\mathbf{P}[A_i] = \left(1 - \frac{1}{n}\right)^m$$

and applying the union bound,

$$\mathbf{P}[A] = \mathbf{P}[\cup_{i=1}^{n} A_i] \le \sum_{i=1}^{n} \mathbf{P}[A_i] = \sum_{i=1}^{n} \left(1 - \frac{1}{n}\right)^m .$$

Suppose we set $m = n \ln n + cn$ for some c. Then

$$\begin{aligned}
\mathbf{P}[A] &\le \sum_{i=1}^{n} \left(1 - \frac{1}{n}\right)^{n(\ln n + c)} \\
&\le \sum_{i=1}^{n} (e^{-1})^{\ln n + c} \\
&= \sum_{i=1}^{n} \frac{1}{ne^c} = \frac{1}{e^c}.
\end{aligned}$$

Thus the probability of failure after $n \ln n + cn$ days is less than $\frac{1}{e^c}$.

Appendix B. Analysis For One Counter Protocol, With Non-Counters Feeling Lucky

Under the one counter protocol, the prisoners escape if and only if the bulb, which is initially OFF, alternates its state from OFF to ON exactly $n-1$ times.

Non-counters can also witness transitions in bulb state, and keep their own tally of prisoners who have been in the room. For example, suppose the first six times a non-counter is interrogated, he sees the sequence

Day	0	15	27	55	95	221	521
Bulb	OFF	ON	ON	OFF, turn bulb ON	ON	OFF	ON

.

Day 0 represents the day before the game begins; although no prisoner is interrogated on this day, all prisoners know that the bulb's initial state is OFF. Given these observations, the non-counter can be sure that at least three distinct non-counters including himself were interrogated between Day 0 and Day 521, since at least three OFF to ON transitions occurred in this time period. He is also sure that the counter was interrogated, since at least one ON to OFF transition occurred (such as from Day 27 to Day 55). Thus, the non-counter can increment his own count to 4.

Thus, if any non-counter witnesses $n-1$ OFF to ON transitions, then that non-counter is equally qualified to declare victory, preempting the standard one counter algorithm. We will call a non-counter is *lucky* if his count reaches $n-1$.[1]

The effects of this improvement on the one counter protocol's performance are negligible. However, for the sake of completeness, an explicit analysis is provided.

B.1. **Expected Runtime Analysis.**

We again use the following nomenclature:

- Let Y_i be the number of days between the $(i-1)$th time the counter turns OFF the bulb, and the ith time a non-counter turns ON the bulb.
- Let Z_i be the number of days between the ith activation of the bulb and the counter tallying it.

Recall that $Y_i \sim \textbf{geom}\left(\frac{n-i}{n}\right)$ and $Z_i \sim \textbf{geom}\,(1/n)$.

Since non-counters can only possibly declare victory after the $(n-2)$th activation of the bulb, the maximum number of days that the proposed improvement can save is precisely $Z_{n-2} + Y_{n-1} + Z_{n-1}$, assuming that the lucky non-counter witnesses the bulb the day immediately after it is turned ON for the last time. In terms of expectation,

$$\begin{aligned}\mathbf{E}\left[X^{(\text{one counter})}\right] - \mathbf{E}\left[X^{(\text{one counter, with non-counters feeling lucky})}\right] &\leq \mathbf{E}\left[Z_{n-1} + Y_{n-1} + Z_{n-1}\right] \\ &= 2\mathbf{E}\left[Z_1\right] + \mathbf{E}\left[Y_{n-1}\right] \\ &= 3n.\end{aligned}$$

[1]The reader may be wondering why the threshold is not $n-2$. Later we prove that a lucky non-counter cannot possibly be the last non-counter to turn ON the bulb.

Thus, this improvement does not affect the average big-Oh runtime of $O(n^2)$ for the One Counter Protocol.

B.2. **Probability Analysis.** We now estimate the probability of a non-counter declaring victory, which we will show to be very, very small.

Before proceeding further, we argue that a lucky non-counter must be either the first or second non-counter to turn ON the light bulb, as follows:

(1) The bulb state sequence is OFF, ON, OFF, ON, ... where the initial OFF is the default bulb state on Day 0.
(2) To reach the maximum count of $n-2$, a lucky non-counter must witness every state of the bulb, except possibly the initial default OFF state (which he can see in his mind).
(3) However, whenever a non-counter sees an OFF bulb for the first time, he must turn ON the bulb.
(4) Thus if a non-counter is not the first or second non-counter to turn ON the bulb, then he could not have seen the OFF bulb resulting from the counter's first deactivation. So such a non-counter must be unlucky.

Rather than saying "first or second non-counter to turn ON the bulb", we will henceforth refer to them as "non-counter with index 1" and "non-counter with index 2".

Let N be the event that there exists a lucky non-counter (who declares victory). Let N_j be the event that the non-counter with index j becomes lucky if the one counter protocol is never preempted. Then by the union bound,

$$\mathbf{P}[N] = \mathbf{P}[N_1 \cup N_2] \leq \mathbf{P}[N_1] + \mathbf{P}[N_2]. \tag{13}$$

We first compute $\mathbf{P}[N_1]$. Let $W_i^{OFF;1} = W_i^{OFF}$ be the event that the non-counter with index 1 witnesses the ith OFF state of the bulb, and let $W_i^{ON;1} := W_i^{ON}$ the event that this non-counter witnesses the ith ON state of the bulb. Then

$$\begin{aligned}\mathbf{P}[N_1] &= \mathbf{P}\left[\bigcap_{i=1}^{n-1} W_i^{OFF} \cap W_i^{ON}\right] \\ &= \underbrace{\mathbf{P}\left[W_1^{OFF}\right]}_{1}\underbrace{\mathbf{P}\left[W_1^{ON}\right]}_{1}\prod_{i=2}^{n-1}\mathbf{P}\left[W_i^{OFF}\right]\mathbf{P}\left[W_i^{ON}\right] \qquad (14) \\ &= \prod_{i=2}^{n-1}\mathbf{P}\left[W_i^{OFF}\right]\mathbf{P}\left[W_i^{ON}\right].\end{aligned}$$

In the second equality, independence follows from the fact that the prisoners are always chosen uniformly at random (u.a.r.). In the third equality, $\mathbf{P}\left[W_1^{OFF}\right] = 1$ since the non-counter knows is initially OFF, and $\mathbf{P}\left[W_1^{ON}\right] = 1$ since the non-counter with index 1 is the first to turn ON the bulb, so he witnesses this transition.

We first address $\mathbf{P}\left[W_i^{ON}\right]$. This probability can be reworded as the probability that the lucky non-counter is selected during the time window when we are waiting for the counter to turn off the bulb for the ith time. ("W" stands for "window".) This time window has random length Z_i. Since all the Z_i's are identically distributed as $\mathbf{geom}\,(1/n)$, $\mathbf{P}\left[W_i^{ON}\right] = \mathbf{P}\left[W_1^{ON}\right]$, and we may write

$$\mathbf{P}\left[N_1\right] = \left(\mathbf{P}\left[W_1^{ON}\right]\right)^{n-2} \prod_{i=2}^{n-1} \mathbf{P}\left[W_i^{OFF}\right].$$

To compute $\mathbf{P}\left[W_1^{ON}\right]$, we condition on Z_1, the day that the counter arrives. The lucky non-counter must arrive to see the ON bulb *before* this day.

$$\begin{aligned}
\mathbf{P}\left[W_1^{ON}\right] &\overset{(a)}{=} \sum_{z=1}^{\infty} \mathbf{P}\left[Z_1 = z\right] \mathbf{P}\left[W_1^{ON} | Z_1 = z\right] \\
&\overset{(b)}{=} \mathbf{P}\left[Z_1 = 1\right] \cdot 0 + \sum_{z=2}^{\infty} \mathbf{P}\left[Z_1 = z\right] \mathbf{P}\left[W_1^{ON} | Z_1 = z\right] \\
&\overset{(c)}{=} \sum_{z=2}^{\infty} \underbrace{\left(\frac{n-1}{n}\right)^{z-1} \left(\frac{1}{n}\right)}_{\mathbf{P}[Z_1 = z]} \underbrace{\left(1 - \left(\frac{n-2}{n}\right)^{z-1}\right)}_{\mathbf{P}\left[W_1^{ON} | Z_1 = z\right]} \qquad (15) \\
&\overset{(d)}{=} \frac{2(n-1)}{3n-2}.
\end{aligned}$$

Each step is justified as follows:

(a) total probability theorem
(b) If the lucky counter arrives immediately after the bulb is switched ON, there is no chance for the lucky non-counter to arrive before the counter does.
(c) $Z_1 \sim \mathbf{geom}\,(1/n)$.
(d) To compute $\mathbf{P}\left[W_1^{ON} | Z_1 = z\right]$, the probability the lucky non-counter arrives before the counter is the complement of the probability that our particular non-counter is not selected in any of the days up to the day the counter arrives.
(e) Follows from the closed-form formula for geometric series and simplifying.

We now compute $\mathbf{P}\left[W_i^{OFF}\right]$ where $i \geq 2$. Let Y_i be the number of days between when the bulb is turned OFF for the ith time and when the bulb is turned ON for the $(i+1)$th time. Conditioning on the value of Y_i and doing a similar analysis,

$$
\begin{aligned}
\mathbf{P}\left[W_i^{OFF}\right] &= \sum_{y=1}^{\infty} \mathbf{P}\left[Y_i = y\right] \mathbf{P}\left[W_i^{OFF} | Y_i = y\right] \\
&\overset{(a)}{=} \mathbf{P}\left[Y_i = 1\right] \cdot 0 + \sum_{y=2}^{\infty} \mathbf{P}\left[Y_i = y\right] \mathbf{P}\left[W_i^{OFF} | Y_i = y\right] \\
&\overset{(b)}{=} \sum_{y=2}^{\infty} \underbrace{\left(\frac{i}{n}\right)^{y-1} \left(\frac{n-i}{n}\right)}_{\mathbf{P}[Y_i = y]} \mathbf{P}\left[W_i^{OFF} | Y_i = y\right] \\
&\overset{(c)}{=} \sum_{y=2}^{\infty} \underbrace{\left(\frac{i}{n}\right)^{y-1} \left(\frac{n-i}{n}\right)}_{\mathbf{P}[Y_i = y]} \underbrace{\left(1 - \left(\frac{i-1}{n}\right)^{y-1}\right)}_{\mathbf{P}\left[W_i^{OFF} | Y_i = y\right]} \\
&\overset{(d)}{=} \frac{i(n-i+1)}{n^2 - i^2 + i}
\end{aligned}
\tag{16}
$$

where each step is justified as follows:

(a) If the bulb is turned ON immediately after it was turned OFF, then there is no opportunity for the lucky non-counter to witness this OFF bulb. The lucky non-counter also could not have turned ON the bulb since he has already done so.
(b) Y_i is a geometric random variable with parameter $\frac{n-i}{n}$.
(c) For the lucky non-counter to see the OFF bulb, he has to chosen on at least one of the $y-1$ days in the OFF time window . We compute the complementary probability that the lucky non-counter is never chosen. Since i non-counters have already activated the bulb, the probability that only the other $i-1$ non-counters (excluding the lucky one) are selected is $\left(\frac{i-1}{n}\right)^{y-1}$.
(d) Follows from the closed-form formula for geometric series.

Thus, combining Equations 15 and 16,

$$
\mathbf{P}\left[N_1\right] = \left(\frac{2(n-1)}{3n-2}\right)^{n-2} \prod_{i=2}^{n-1} \frac{i(n-i+1)}{n^2 - i^2 + i}. \tag{17}
$$

The calculation of $\mathbf{P}[N_2]$ is almost identical, except for a few initial conditions:

$$\begin{aligned}
\mathbf{P}[N_2] &= \mathbf{P}\left[\bigcap_{i=1}^{n-1} W_i^{OFF;2} \cap W_i^{ON;2}\right] \\
&= \underbrace{\mathbf{P}\left[W_1^{OFF}\right]}_{1} \mathbf{P}\left[W_1^{ON}\right] \underbrace{\mathbf{P}\left[W_2^{OFF}\right]}_{1} \underbrace{\mathbf{P}\left[W_2^{ON}\right]}_{1} \prod_{i=3}^{n-1} \mathbf{P}\left[W_i^{OFF}\right] \mathbf{P}\left[W_i^{ON}\right] \\
&\overset{(a)}{=} \mathbf{P}\left[W_1^{ON}\right] \prod_{i=3}^{n-1} \mathbf{P}\left[W_i^{OFF}\right] \mathbf{P}\left[W_i^{ON}\right] \\
&= \mathbf{P}\left[W_1^{ON}\right]^{n-2} \prod_{i=3}^{n-1} \mathbf{P}\left[W_i^{OFF}\right] \\
&\overset{(b)}{=} \left(\frac{2(n-1)}{3n-2}\right)^{n-2} \prod_{i=3}^{n-1} \frac{i(n-i+1)}{n^2-i^2+i}
\end{aligned} \tag{18}$$

where each step is justified as follows:

(a) $\mathbf{P}\left[W_1^{OFF}\right] = 1$ since all prisoners know that the bulb is initially OFF. $\mathbf{P}\left[W_2^{OFF}\right]$ and $\mathbf{P}\left[W_2^{ON}\right]$ are both 1 since the lucky non-counter now has index 2, and is responsible for flipping the bulb ON for the second time.

(b) Substitute Equations 15 and 16.

Finally, plugging these results into Equation 13,

$$\begin{aligned}
\mathbf{P}[N] &\leq \mathbf{P}[N_1] + \mathbf{P}[N_2] \\
&= \left(\frac{2(n-1)}{3n-2}\right)^{n-2} \left(\prod_{i=3}^{n-1} \frac{i(n-i+1)}{n^2-i^2+i}\right) \left(\frac{n^2+2n-4}{n^2-2}\right)
\end{aligned} \tag{19}$$

This could be rewritten in terms of Gamma functions, but we will not bother. Here is a table of values which shows how quickly this upper bound on the probability of a non-counter declaring victory vanishes.

n	Upper bound on $\mathbf{P}[N]$
3	0.898
4	0.309
5	0.0916
6	0.0247
7	0.00622
8	0.00149
9	0.000342
10	0.0000763
15	3.06×10^{-8}
20	8.94×10^{-12}
25	2.19×10^{-15}
50	5.84×10^{-34}
75	6.81×10^{-53}
100	5.63×10^{-72}
1000	2.53×10^{-769}

References

1. P-O. Dehaye, D. Ford, and H. Segerman, *One hundred prisoners and a lightbulb*, Mathematical Intelligencer **25(4)** (2003), 53–61.
2. IBM Research, *100 prisoners and a lightbulb challenge*, `http://domino.watson.ibm.com/Comm/wwwr_ponder.nsf/challenges/July2002.html`, 2002.
3. W. Wu, *[wu:forums]*, William Wu's discussion forum of mathematical riddles and puzzles. `http://wuriddles.com`, `http://www.ocf.berkeley.edu/~wwu/cgi-bin/yabb/YaBB.cgi`, 2002.

Information Systems Lab, Stanford University

E-mail address: `willywu@stanford.edu`