

A Line of Sages

Tanya Khovanova
MIT

June 22, 2013

Abstract

A new variation of an old hat puzzle, where sages are standing in line one behind the other.

For many years this was my favorite hat puzzle:

A sultan decides to give 100 of his sages a test. He has the sages stand in line, one behind the other, so that the last person in line can see everyone else. The sultan puts either a black or a white hat on each sage. The sages can only see the colors of the hats on all the people in front of them. Then, in any order they want, each sage guesses the color of the hat on his own head. Each hears all previously made guesses, but other than that, the sages cannot speak. Each person who guesses the color wrong will have his head chopped off. The ones who guess correctly go free. The rules of the test are given to them one day before the test, at which point they have a chance to agree on a strategy that will minimize the number of people who die during this test. What should that strategy be?

You may now pause and try to solve this puzzle, as the solution is coming below. Also, there will be more puzzles in this essay, and I'll provide the solutions to those as well. It's up to you to decide whether to solve each puzzle by yourself as you read it, or to continue reading on to the solutions. I will inject at least one paragraph of text between each problem and its solution, which gives you several seconds to come up with your own ideas before getting to the solution.

Back to the puzzle. The sages are not allowed to communicate. That means they can't cough or sneeze or poke the person in front of them. They can't even change the pitch of their voices. They can only say either "white" or "black."

Many people expect that about half of the sages will die. Indeed, if they were to guess randomly that would be the case. Surprisingly, only one person actually risks his life. Clearly, the last person has no way of knowing the color of his hat. So it is he who may sacrifice his life to save all the others.

Here is the strategy. The answer of the sage who is last in line signals the parity of the number of black hats he sees. In other words, if we assign zero to a white hat and one to a black hat, the last sage broadcasts "white" if he sees an even number of black hats ahead

of him, and “black” if he sees an odd number of black hats. Then the sages declare their colors in the order from back to front. When it is time for each sage to speak, after having heard all the people behind him and having seen all the people in front of him, he announces “white” if he heard and saw an even number of blacks and “black” otherwise.

A natural next puzzle: What happens if we have many colors? What would be the strategy in this case?

These problems for two and three colors first appeared in 23-rd All-Russian Mathematical Olympiad in 1997 and were submitted by Konstantin Knop [1].

We assume that the sages have been shown the hats and know the colors before the test begins. Actually, we do not need to assume that. We can assume that somewhere there is a list of N colors and that the list is known; the hat colors have to be from this list. We also need to assume that the sages can distinguish colors from any distance. By the way, I am not a sage: I can’t distinguish aqua from turquoise at any distance, but luckily I am not in that line.

Surprisingly, the number of colors does not increase the number of sages risking their lives.

So the strategy is the following: We assign a number to each color and the last person says the number, that is the color, of the sum of the colors he sees modulo the total number of colors, N .

This puzzle is based on error-correcting codes. Suppose we attach a digit to a number so that the sum of digits of the resulting number is divisible by 10. Then, if during a transmission of this number one of the digits gets garbled and becomes unreadable, we can recover it. Similarly, the second to last sage knows the sum of all colors modulo N and sees all the colors but one, so he can calculate his color. Proceeding from back to front, all the other sages can calculate their colors.

I recently received a beautiful variation of this puzzle by Konstantin Knop and Alexander Shapovalov:

A sultan decides to give 100 of his sages a test. The sages will stand in line, one behind the other, so that the last person in the line sees everyone else. The sultan has 101 hats, each of a different color, and the sages know all the colors. The sultan puts all but one of the hats on the sages. The sages can only see the colors of the hats on people in front of them. Then, in any order they want, each sage guesses the color of the hat on his own head. Each hears all previously made guesses, but other than that, the sages cannot speak. They are not allowed to repeat a color that was already announced. Each person who guesses his color wrong will get his head chopped off. The ones who guess correctly go free. The rules of the test are given to them one day before the test, at which point they have a chance to agree on a strategy that will minimize the number of people who die during this test. What should that strategy be?

One might say that this is the same puzzle as the preceding puzzle. Just that every hat is of a different color. The last person will say the sum of the colors of the hats in front of

him modulo 101, and from back to front the other sages will calculate the colors of their own hats. Unfortunately, the sages are not allowed to repeat colors. If one person says his hat is red, then another person can't say that his hat is red even if it is.

Let's see if the previous method can be adjusted to this new annoying constraint. The last sage says the sum of the colors he sees modulo 101. With probability 99/101 this is one of the colors that is used by people in front of him. So we are almost guaranteed that there is an unlucky sage who has the same hat color that was announced by the last person. The sages start calculating their colors from back to front and declaring them until the line reaches the doomed guy who doesn't know what to do. He knows his color, but he isn't allowed to say it.

What does that mean—not allowed? Assume for a second that the sultan is not as cruel as he really is and that the punishment for repeating a color is the death of the offender only. In this case the unlucky man has to repeat his color to avoid compromising the strategy. He saves everyone else, but he will die, as will the last sage. But in real life, the sultan is cruel: if someone repeats a color, everyone is dead. Oops.

Can this strategic idea be saved? If a sage has the same color as the one that was announced by the last person, is there a way to avoid repeating the color and sacrificing many people? If the ill-fated sage says any other color that wasn't used, he will throw off the strategy and compromise everyone in front of him. Is there a way to indicate that he is replacing his color by a bogus color?

The authors of the problem suggest the following idea: The unfortunate sage can say the color of the hat of the first person in line. Everyone but the first person sees this color. All the sages who didn't proclaim their colors yet would know that if someone other than the last sage announces the color of the first hat, then this person is the unlucky one who has the hat of the color stated by the sage who is last in line. So people in front can adjust their calculations and proceed guessing their colors correctly. That is, until the first person is reached, for he is doomed. In this solution, not more than three people sacrifice their lives to save everyone else.

Actually, even though three people might die, only one of them is a hero. The last person knows his color with probability one half. By announcing one of the colors he sees, he forgoes his chance of survival and sacrifices himself by helping others. The other two guys are dead no matter what—their colors are already used before their turn.

This problem appeared at the Tournament of the Towns in March 2013 [2]. It had two difficulty levels; the first one asked to save at least half of the people, which we over-achieved. The second part asked if it is possible to save all but one.

There is more to this puzzle than we just discussed. It is indeed possible to save all but one sage. Then why on earth did I discuss the non-optimal solution? First, it is pretty. Second, the solution generalizes to the situation when the sultan has 102 or more hats of different colors.

But we need to get back to this puzzle and figure out how to save all but one. It is the last sage who should be sacrificed, because there is never enough information to guarantee his survival. If we want to save everyone else, the last person can't announce a color that he sees. That is, he has a choice of two colors: the one on his head and the color of the unused hat. How can he signal to everyone what to say by picking one of the two colors?

It turns out that this problem is not about number theory and moduli. This puzzle is

about permutations. We can imagine that there is a ghost sage behind the last sage and that he has the leftover hat. We can assign numbers to colors, and the arrangement of hats then becomes a permutation of length 101.

So, what can the last real (non-ghost) sage signal with one bit of information? He chooses the color out of the two he doesn't see so that the permutation is even. After that the sages unravel their colors from back to front. Each has a choice between two colors he doesn't see or hasn't heard another sage say. Essentially, each person chooses between two permutations of length 101 that differ by one transposition. So each sage picks the even permutation. This way everyone but the last person goes free. And the last person doesn't need to be a hero. He survives with a probability of one half.

Now I am not sure which hat puzzle is my favorite: the classic one or the new one.

References

- [1] 23-rd All-Russian Mathematical Olympiad 1997, <http://www.imomath.com/othercomp/Rus/RusM097.pdf>, 1997.
- [2] Thirty Fourth Tournament of the Towns, Spring, <http://www.turgor.ru/problems/34/vs34s1-avt.pdf> (in Russian), 2013.