An Analysis of the Game Tenzi

Justin Barhite
Carthage College
jbarhite@carthage.edu

June 15, 2016

Abstract

In the game Tenzi, a player has ten six-sided dice and repeatedly rolls any subsets of the dice in order to have all ten show the same number with as few rolls as possible. A simple strategy is to choose a number at the start of the game and always reroll any dice not showing that number. For this strategy we find that the expected number of rolls to win the game is 16.6, and we derive the expected number of rolls with playing the game with any number of dice instead of ten.

1 Introduction

In the game Tenzi, each player has ten dice and the goal of getting all ten to show the same number. The player begins by rolling all ten dice, then rerolls any subset of the dice, and continues rerolling until all ten dice are the same. On each roll, the player is free to choose whichever and however many dice to reroll. Multiple people playing a game of Tenzi roll and reroll at will, without taking turns or waiting for other players, and the first to get ten matching dice wins the game.

Several questions naturally arise regarding the mathematical structure of the game. How many rolls, on average, does it take for a player to get all ten dice to be the same? How is that average number of rolls affected by the strategy employed? What is the optimal strategy to minimize the number of rolls? How do these answers change when the game is played with N dice instead of 10? By modeling Tenzi as a Markov chain, we derive the expected number of rolls to win for one particular strategy and show that it behaves as expected in the large-N limit. We can also make reasonable guesses at the answers to the other questions based on computational results.

2 Background

Given any possible state of the dice, a strategy will determine which dice will be rerolled. There are two strategies we consider, which we call the choose-a-number strategy and the largest-group strategy. The choose-a-number strategy is so named because at the beginning of the game, before the first roll of the dice, we select a number from 1 to 6; that becomes the number we attempt to collect as the game progresses. At each step we keep any dice showing the chosen number and reroll all the others. However, we could do better by instead keeping the largest single group of dice showing the same number (breaking ties arbitrarily) and rerolling all others; this is the largest-group strategy. The largest-group strategy intuitively seems better, and computational results show that it is, but the choose-a-number strategy is almost as good, and much easier to analyze mathematically.

We model a game of Tenzi with a Markov chain. A Markov chain comprises a set of discrete states between which the system transitions, supplemented by the Markov property, which requires that the probability of transitioning to a given state depend only on the current state, not on any history of states. After each roll of the dice in a game of Tenzi, we will represent the game with one of eleven states, numbered from 0 to 10, which describe our progress toward winning. If we play the choose-a-number strategy, the state will be the number of correct dice (i.e. dice showing the chosen number) we have; with the largest-group strategy the state is the size of the largest group. In either case, the game begins in state 0 before the dice are first rolled, and the game ends once it transitions to state 10. It is clear that for these two strategies, and presumably any other reasonable strategy, the probability of transitioning to another state depends only on the strategy and the current state of the dice.

Once we choose a strategy, we can compute the probabilities of transitioning from any state to any other. We write these probabilities in a transition matrix, where each row and column represents a state and each entry gives the probability of transitioning from the row state to the column state. For instance, if we adopt the choose-a-number strategy and play with two dice instead of ten, there are only three states instead of eleven, and the transition matrix is[1]

$$\mathbf{M} = \begin{bmatrix} \left(\frac{5}{6}\right)^2 & 2\left(\frac{5}{6}\right)\left(\frac{1}{6}\right) & \left(\frac{1}{6}\right)^2 \\ 0 & \left(\frac{5}{6}\right) & \left(\frac{1}{6}\right) \\ 0 & 0 & 1 \end{bmatrix}. \tag{1}$$

The first row gives the probabilities of transitioning from the initial state of zero correct dice to any other state; for example, $M_{12} = 2(5/6)(1/6)$ is the probability that when the two dice are rolled, one of them shows the chosen number and the other does not. The second row gives transition probabilities from the state with one correct die; we note that $M_{21} = 0$ since the one correct die will not be rerolled, so there is no chance of returning to the state of zero correct dice. Finally, the state of two correct dice is known as an absorbing

[1]This actually assumes that we only win when both dice show the chosen number, whereas under the normal rules we would still win if, for example, we chose 1 but immediately rolled all 3s. This distinction has negligible effect if we play with more than a few dice.

state; once this state is reached the game will not transition to any other state (the game has been won).

3 Expected Number of Rolls

The expected number of rolls $E(N)$ to win the game with N dice is

$$E(N) = \sum_{k=0}^{\infty} \sum_{j=1}^{N} (\mathbf{M}^k)_{1,j}. \tag{2}$$

We will compute the expected number of rolls for the choose-a-number strategy, which works because it is a simple matter to write out the transition matrix for a given N; the largest-group strategy is substantially more difficult to analyze since the transition matrix for arbitrary N is more difficult to write in closed form.

3.1 Raising the Transition Matrix to a Power

We will first consider the two-dice case and then generalize to N dice. The transition matrix for two dice is

$$\mathbf{M} = \begin{bmatrix} a^2 & 2ab & b^2 \\ 0 & a & b \\ 0 & 0 & 1 \end{bmatrix}, \tag{3}$$

where for convenience we define $a := 5/6$ and $b := 1/6$.[2] Now suppose that

$$\mathbf{M}^k = \begin{bmatrix} a^{2k} & 2a^k b \sum_{l=0}^{k-1} a^l & \left(b \sum_{l=0}^{k-1} a^l \right)^2 \\ 0 & a^k & b \sum_{l=0}^{k-1} a^l \\ 0 & 0 & 1 \end{bmatrix} \tag{4}$$

for $k \in \mathbb{N}$, and note that this reduces to Eq. (3) when $k = 1$. Multiplying $\mathbf{M}^k$ and $\mathbf{M}$ we find that

$$\mathbf{M}^{k+1} = \begin{bmatrix} a^{2(k+1)} & 2a^{k+1} b \sum_{l=0}^{k} a^l & \left(b \sum_{l=0}^{k} a^l \right)^2 \\ 0 & a^{k+1} & b \sum_{l=0}^{k} a^l \\ 0 & 0 & 1 \end{bmatrix}, \tag{5}$$

so by induction we conclude that Eq. (4) holds for all $k \in \mathbb{N}$. For N dice, the transition matrix has dimension $(N+1) \times (N+1)$ and has the form

$$\mathbf{M}_N = \left[\begin{array}{c|cccc} a^N & \binom{N}{1} a^{N-1} b & \binom{N}{2} a^{N-2} b^2 & \cdots & b^N \\ \hline \mathbf{0} & & \mathbf{M}_{N-1} & & \end{array} \right], \tag{6}$$

[2] Any of the results that follow can be generalized to n-sided dice by instead letting $b := 1/n$ and $a := 1 - b$.

and powers of $\mathbf{M}_N$ are given by

$$\mathbf{M}_N^k = \left[\begin{array}{c|ccccc} a^{Nk} & \binom{N}{1}a^{(N-1)k}bS_k & \binom{N}{2}a^{(N-2)k}(bS_k)^2 & \cdots & (bS_k)^N \\ \hline \mathbf{0} & & & \mathbf{M}_{N-1}^k & \end{array}\right], \tag{7}$$

where we have defined

$$S_k := \sum_{l=0}^{k-1} a^l. \tag{8}$$

As with the two-dice matrix, the expression for $\mathbf{M}_N^k$ is verified inductively by computing $\mathbf{M}_N^{k+1} = \mathbf{M}_N^k \mathbf{M}_N$. To demonstrate this, we will compute $\left(\mathbf{M}_N^{k+1}\right)_{1,j}$ for an arbitrary $j \in \{2, 3, \ldots, N+1\}$; it is fairly easy to see that for the rest of the matrix, if Eq. (7) holds for $\mathbf{M}_N^k$ then it holds for $\mathbf{M}_N^{k+1}$. Multiplying the matrices, we have

$$\left(\mathbf{M}_N^{k+1}\right)_{1,j} = \sum_{i=1}^{N+1} \left(\mathbf{M}_N^k\right)_{1,i} \cdot \left(\mathbf{M}_N\right)_{i,j}. \tag{9}$$

The first factor inside this sum can be read from the first row of the matrix in Eq. (7):

$$\left(\mathbf{M}_N^k\right)_{1,i} = \binom{N}{i-1} a^{(N+1-i)k} \left(bS_k\right)^{i-1}. \tag{10}$$

The second factor is a little harder to determine since we defined the transition matrix recursively. It is easily shown by induction that $\mathbf{M}_N$ is an upper triangular matrix for any N, so $\left(\mathbf{M}_N\right)_{i,j} = 0$ if $i > j$. If $i \leq j$, then we see from Eq. (7) that

$$\left(\mathbf{M}_N\right)_{i,j} = \left(\mathbf{M}_{N-(i-1)}\right)_{1,j-(i-1)} \tag{11}$$

This can be made rigorous by inducting on i. We can now use the first row of the matrix in Eq. (6) to find the exact value of this entry:

$$\left(\mathbf{M}_N\right)_{i,j} = \binom{N-(i-1)}{j-i} a^{N-j+1} b^{j-i}. \tag{12}$$

Substituting Equations (10) and (12) into Eq. (9), we obtain

$$\left(\mathbf{M}_N^{k+1}\right)_{1,j} = \sum_{i=1}^{j} \left[\binom{N}{i-1} a^{(N+1-i)k} \left(bS_k\right)^{i-1} \cdot \binom{N-(i-1)}{j-i} a^{N-j+1} b^{j-i}\right]. \tag{13}$$

Note that we only have to sum from 1 to j since $\left(\mathbf{M}_N\right)_{i,j} = 0$ for $i > j$. Factoring out powers of a and b, we have

$$\left(\mathbf{M}_N^{k+1}\right)_{1,j} = a^{(N-j+1)(k+1)} b^{j-1} \sum_{i=1}^{j} \left[\binom{N}{i-1} \binom{N-(i-1)}{j-i} {S_k}^{i-1} a^{(j-i)k}\right]. \tag{14}$$

We now replace the two binomial coefficients with a product of two different binomial coefficients, and we claim that their equality is verified by writing out all four binomial coefficients in terms of factorials:

$$\left(\mathbf{M}_N^{k+1}\right)_{1,j} = a^{(N-j+1)(k+1)} b^{j-1} \binom{N}{j-1} \sum_{i=1}^{j} \left[\binom{j-1}{i-1} S_k{}^{i-1} \left(a^k\right)^{j-i} \right]. \quad (15)$$

Finally, we invoke the binomial theorem to obtain

$$\left(\mathbf{M}_N^{k+1}\right)_{1,j} = a^{(N-j+1)(k+1)} b^{j-1} \binom{N}{j-1} \left(S_k + a^k\right)^{j-1}, \quad (16)$$

which is the desired result,

$$\left(\mathbf{M}_N^{k+1}\right)_{1,j} = \binom{N}{j-1} a^{(N-j+1)(k+1)} \left(bS_{k+1}\right)^{j-1}. \quad (17)$$

Therefore we conclude that Eq. (7) correctly gives the N-dice transition matrix raised to any power.

3.2 Computing the Expected Number of Rolls

The expected number of rolls is given by Eq. (2). Since the first row of $\mathbf{M}^k$ sums to 1, we can rewrite this as

$$E(N) = \sum_{k=0}^{\infty} \left[1 - (\mathbf{M}^k)_{1,N+1}\right]. \quad (18)$$

We then compute that

$$E(N) = \sum_{k=0}^{\infty} \left[1 - \left(b \sum_{l=0}^{k-1} a^l \right)^N \right]. \quad (19)$$

We can rewrite the geometric sum to obtain

$$E(N) = \sum_{k=0}^{\infty} \left[1 - \left(b \cdot \frac{1-a^k}{1-a} \right)^N \right], \quad (20)$$

and since $a + b = 1$ this reduces to

$$E(N) = \sum_{k=0}^{\infty} \left[1 - \left(1 - a^k\right)^N \right]. \quad (21)$$

This expression will be useful for investigating the behavior of $E(N)$ for large N, but we can recast it as a finite sum, which may be easier to evaluate numerically for some values of N. Applying the binomial theorem, we have

$$E(N) = \sum_{k=0}^{\infty} \left[1 - \sum_{n=0}^{N} \binom{N}{n} \left(-a^k\right)^n \right], \quad (22)$$

which simplifies to

$$E(N) = -\sum_{k=0}^{\infty} \left[\sum_{n=1}^{N} \binom{N}{n} (-1)^n \left(a^k\right)^n \right]. \tag{23}$$

Reordering the sums, we have

$$E(N) = -\sum_{n=1}^{N} \left[\binom{N}{n} (-1)^n \sum_{k=0}^{\infty} (a^n)^k \right], \tag{24}$$

and evaluating the infinite sum we have

$$E(N) = \sum_{n=1}^{N} \left[\binom{N}{n} \frac{(-1)^{n+1}}{1-a^n} \right]. \tag{25}$$

For the $N = 10$ case this evaluates to $E(10) \approx 16.6$.

3.3 Approximating the Number of Rolls

To discern the general form of expected rolls as a function of the number of dice, consider playing the game with $(6/5)N$ dice. If N is large, then on the first roll we expect the dice to be fairly evenly divided between the numbers 1 through 6. Accordingly, we expect that about $1/6$ of those dice will be correct, leaving roughly the other $5/6 \cdot (6/5 \cdot N) = N$ to be rerolled on the second roll. Since the first roll of the game reduces the number of dice being rolled from $(6/5)N$ to N, we might conjecture that

$$E(6/5 \cdot N) = E(N) + 1, \tag{26}$$

and this sort of relationship suggests that the general form of the expected number of rolls is

$$E(N) \approx \log_{6/5} N. \tag{27}$$

To verify this by a more direct approach, we approximate the sum in Eq. (21) as an integral,

$$E(N) \approx \int_0^{\infty} \left[1 - (1-a^x)^N\right] \mathrm{d}x \tag{28}$$

Making the substitution $u = 1 - a^x$, this simplifies to

$$E(N) \approx \int_0^1 \left(1 - u^N\right) \cdot \frac{-\mathrm{d}u}{(1-u)\log a} \tag{29}$$

$$= \frac{1}{\log(1/a)} \int_0^1 \frac{1-u^N}{1-u} \mathrm{d}u. \tag{30}$$

The integral evaluates to the Nth harmonic number H_N, and since the harmonic numbers increase approximately as $\log N$, we have the expected result,

$$E(N) \approx \frac{H_N}{\log 6/5} \approx \log_{6/5} N. \tag{31}$$

4 Conclusion

We have computed the expected number of rolls to win a generalized game of Tenzi using N dice and playing the choose-a-number strategy. It remains to justify approximating the sum in Eq. (21) by an integral, but assuming this approximation is valid we have shown the expected number of rolls to grow logarithmically in N, as expected. Computational results show that the largest-group strategy is slightly better, with an expected 15.3 rolls to win using 10 dice (as compared to 16.6 for choose-a-number), but the difference between the two strategies is observed to decrease as N becomes large; future work will investigate whether the difference converges to zero. We have not found any better strategy than the largest-group strategy, but is still to be seen whether the largest-group strategy is optimal or a better one can be found.