

JUGGLING PROBABILITIES

GREGORY S. WARRINGTON

1. INTRODUCTION

Imagine¹ yourself effortlessly juggling five balls in a high, lazy pattern. Your right hand catches a ball and immediately throws it. One second later, it is your left hand's turn to catch and throw a ball. Then it is your right hand's turn again. . . . Some of your throws may be low; some high. Some balls go straight up; others cross over to the opposite hand. At most times, one ball lands; occasionally a hand remains empty. But the alternating cadence of your hands is unwavering.

Suppose that, while you are juggling, we momentarily pause time. Certain balls are in the air — you have already thrown them. To avoid dropping the ball, you must make catches at certain times in the future. If your previous few throws were all low, perhaps you are only committed at 1,2,4,6 and 7 seconds in the future. On the other hand, if you had just vigorously launched a ball, you might be committed at 1,2,3,4 and 10 seconds in the future. The set of “committed times” is your *landing state* (“state” for short). As you juggle, you wander from state to state according to what throw you have most recently made.

Our goal in this paper is to answer the following

Random Question. What fraction of the time is spent in any given state if every throw is chosen randomly?

The answer, of course, depends on how we specify the following parameters:

1. the possible states,
2. the legal throws from a given state and
3. the probability of making each legal throw.

There are countless ways to make the above specifications, but we will begin our investigations with the model that most closely mirrors what people envision as juggling. In the next section, we describe the juggling universe determined by this standard model and answer our Random Question. In Section 3 we generalize the notion of a state in order to give a simple proof of our answer. Finally, in Section 4, we describe variations of our model.

2. JUGGLING STATES

Jugglers have been known to {toss, roll, spin, bounce, drop, kick}
 their {balls, rings, clubs, flames, chairs}
 while accompanying their act with {patter, music, dance, flourishes}.

We, however, will strip juggling down to what is (arguably) its essentials. Our juggler will juggle only balls and only in the air. He will not attempt jokes. Whether a hand is under a leg or behind the back when making a catch will not be noted. *We only record the order in which the balls are thrown and caught.* While clearly ignoring much of what makes juggling visually intriguing, our single-mindedness will free us to highlight the inherently mathematical nature of juggling.

Our juggler will:

1. throw alternately from each hand, once a second,
2. throw a ball instantaneously upon catching it,

Date: July 18, 2010.

¹The author would like to thank Peter Doyle and J. Laurie Snell for their substantial help in simplifying the original proof of the main result of this paper.

3. be assumed to have always been (and be forevermore) juggling
(this allows us to avoid starting/stopping issues),
4. be named Magnus².

Every ball Magnus throws lands some (integral) number of seconds in the future. By abuse of physics, we will refer to this parameter as the *height* of the throw. Until mentioned otherwise, we now fix some maximum height h above which Magnus is too weak to throw.

In the introduction, we referred to a state as a set of future times at which balls are landing. It will be more convenient throughout the paper to define a (*landing*) *state* ν as an h -tuple in the alphabet $\{\bullet, \smile\}$. Given a state ν , we will write ν_t for its t -th component and write ν as

$$\nu = \boxed{\nu_1 \nu_2 \cdots \nu_h}.$$

If $\nu_t = \bullet$, then Magnus will catch a ball t seconds in the future. If $\nu_t = \smile$, then no ball lands t seconds in the future. We make that convention that $\nu_t = \smile$ for $t > h$. To recover our original notion of a state from ν , take the set $\{t : \nu_t = \bullet\}$. We will denote the set of states that are h -tuples by St_h . The subset $St_{h,f} \subset St_h$ will contain only those states with f \smile 's.

As an example, consider the state $\nu = \boxed{\bullet \bullet \smile \bullet \smile \bullet \bullet \smile \smile} \in St_{10}$. Assuming that Magnus has just thrown out of his left hand, at 1 and 7 seconds in the future he will be (instantaneously) throwing/catching out of his right hand; at 2,4,6 and 8 seconds in future he will be (instantaneously) throwing/catching out of his left hand.

When the first element of our state ν is \smile , we will often refer to what Magnus does during the next second as making a “height 0” throw. Upon making a throw (even a 0), Magnus commits himself to a (usually different) new state ω :

1. If the throw is a 0, $\omega = \boxed{\nu_2 \nu_3 \cdots \nu_h \smile}$.
2. If the throw is of height $t \geq 1$, then $\omega = \boxed{\nu_2 \nu_3 \cdots \bullet \cdots \nu_h \smile}$, where the \bullet is placed in the t -th position (replacing ν_{t+1}).

As an example, below we illustrate successive throws of height 2, 0 and 6 from an initial state in which balls are landing at 1, 4 and 6 seconds in the future:

$$\boxed{\bullet \smile \smile \bullet \bullet} \xrightarrow{2} \boxed{\smile \bullet \bullet \smile \smile} \xrightarrow{0} \boxed{\bullet \bullet \smile \bullet \smile} \xrightarrow{6} \boxed{\bullet \smile \smile \smile \bullet}.$$

In Figure 1, we show all possible states and transitions when Magnus has 3 balls and is not able to throw anything higher than a 4.

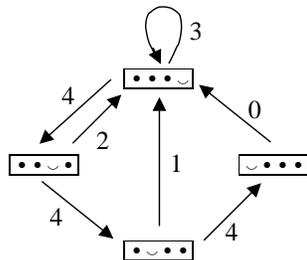


FIGURE 1. Juggling 3 balls with maximum height 4. The edges are labeled with the throw height.

Remark 1. In this section, we have introduced the fundamentals of “siteswap” notation — albeit in an untraditional form. As with many great ideas, the provenance of siteswap notation is somewhat murky. However, we can safely say that this notation was introduced (independently, and in various forms) in the mid 1980s by Paul Klimek; Colin Wright and members of the Cambridge Juggling Club; and Bengt Magnusson and Bruce Tiemann. For a more leisurely introduction and a bibliography of papers relating to

²After Magnus Nicholls, a juggler prominent during the early 1900s. He is reputed to be the first person to juggle five clubs. Reliable details of his life are outclassed by fascinating stories.

siteswap notation, the reader should consult [4]. General internet resources on juggling include [1] and [2]. A recent book devoted to the mathematics of juggling is [8].

Both as a useful notation for juggling and for interesting enumerative combinatorics, it is preferable to introduce siteswap notation by defining patterns as repeating sequences of throws. Such patterns correspond to cycles in certain graphs; such a graph can be found in Figure 1. Many computer programs have been written to animate these patterns. See [3] for one cowritten by the author, or [2] for a comprehensive list.

To answer our Random Question, we need an explicit description of the set of possible states Magnus can pass through and of the possible transitions among them. To describe these transitions, we introduce a family of *throwing* operators Θ_j acting on these states. Intuitively, Θ_0 corresponds to Magnus waiting when no ball lands and Θ_j corresponds to Magnus making a height j throw when a ball does land. Precisely,

$$(1) \quad \Theta_j(\nu) = \begin{cases} \boxed{\nu_2 \cdots \nu_h \smile}, & \text{if } j = 0, \\ \boxed{\nu_2 \cdots \nu_j \bullet \nu_{j+2} \cdots \nu_h \smile}, & \text{if } j > 0. \end{cases}$$

For the remainder of this and the next section, we will fix $h > 0$ and $0 \leq f \leq h$. We are now ready to define the directed graph which specifies the possible states for Magnus along with his possible transitions. Informally, the valid transitions correspond either to waiting or to making throws to empty landing times.

Definition 2. The *state graph* $G_{h,f}$ is the directed graph with

1. Vertices indexed by $St_{h,f}$.
2. An edge from ν to ω whenever
 - (a) $\omega = \Theta_j(\nu)$ for some j with $1 \leq j \leq h$, $\nu_1 = \bullet$ and $\nu_{j+1} = \smile$; or
 - (b) $\omega = \Theta_0(\nu)$ and $\nu_1 = \smile$.

The set of all edges is denoted $\text{Edges}(G_{h,f})$.

If $(\nu, \omega) \in \text{Edges}(G_{h,f})$, then we will refer to ν as a *precursor* of ω .

Figure 2 illustrates $G_{5,2}$.

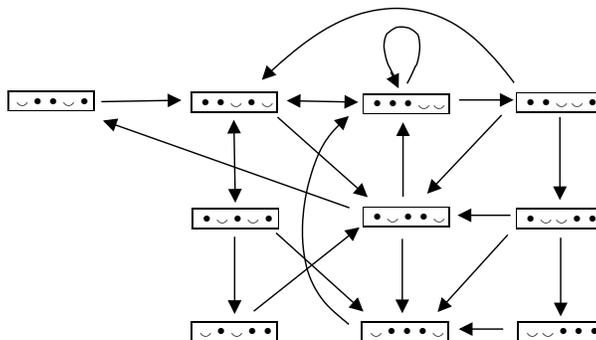
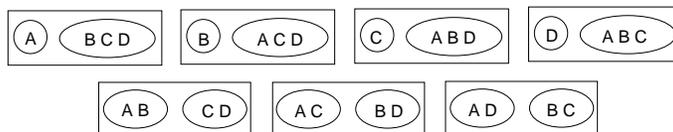


FIGURE 2.

Let us revisit the scenario introduced at the beginning of the paper (this time with Magnus juggling). When we freeze the action, Magnus is in a particular state — he is committed to making certain catches. If no ball is landing in one second, he is forced to wait. If a ball is landing, however, he can throw to any height ($\leq h$) that will not lead to two balls landing at the same time. Our assumption in the following is that Magnus chooses, with equal probability, one of these throws. In this manner, Magnus is taking a random walk on the graph $G_{h,f}$. Having now phrased the question in a precise manner, we need only give two definitions to state the answer.

Definition 3. The *Stirling number of the second kind*, $\left\{ \begin{smallmatrix} a \\ b \end{smallmatrix} \right\}$, counts the number of ways of partitioning an a -element set into b blocks (irrespective of order).

For example, $\left\{ \begin{smallmatrix} 4 \\ 2 \end{smallmatrix} \right\} = 7$ as we can partition $\{A, B, C, D\}$ into two parts in seven ways:



	1	2	3	4	5	6
1	1					
2	1	1				
3	1	3	1			
4	1	7	6	1		
5	1	15	25	10	1	
6	1	31	90	65	15	1

TABLE 1. Stirling numbers of the second kind $\left\{ \begin{smallmatrix} a \\ b \end{smallmatrix} \right\}$.

Let $\nu \in St_h$. For each integer t with $1 \leq t \leq h$, define

$$(2) \quad \phi_t(\nu) = \begin{cases} |\{t < j \leq h : \nu_j = \smile\}|, & \text{if } \nu_t = \bullet, \\ 0, & \text{if } \nu_t = \smile. \end{cases}$$

Definition 4. The *weight* of a state ν , $\Delta(\nu)$, is given by the formula

$$(3) \quad \Delta(\nu) = \prod_{t=1}^h (1 + \phi_t(\nu)).$$

The weight (our terminology) arises when counting permutations with restricted positions. See, e.g., [9, 2.4].

We are finally in a position to answer our Random Question.

Theorem 5. *If Magnus is juggling b balls and is able to throw them to heights at most $h = b + f$, then over the long term he spends*

$$(4) \quad \frac{\Delta(\nu)}{\left\{ \begin{smallmatrix} h+1 \\ f+1 \end{smallmatrix} \right\}}$$

of his time in the state $\nu \in St_{h,f}$.

Example 6. Below we present the vector α whose ν -th component specifies the fraction of the time Magnus spends in each state ν of Figure 2 while juggling 3 balls at a maximum throw height of 5.

$$\begin{aligned} \alpha &= \frac{1}{\left\{ \begin{smallmatrix} 6 \\ 3 \end{smallmatrix} \right\}} (\Delta(\boxed{\bullet \bullet \smile \smile}), \Delta(\boxed{\bullet \smile \bullet \smile}), \Delta(\boxed{\bullet \smile \smile \bullet}), \\ &\quad \Delta(\boxed{\smile \bullet \bullet \smile}), \Delta(\boxed{\smile \bullet \smile \bullet}), \Delta(\boxed{\smile \smile \bullet \bullet}), \Delta(\boxed{\smile \bullet \bullet \smile}), \\ &\quad \Delta(\boxed{\smile \bullet \smile \bullet}), \Delta(\boxed{\smile \smile \bullet \bullet}), \Delta(\boxed{\smile \smile \bullet \bullet})) \\ &= \frac{1}{90} (27, 18, 9, 12, 6, 3, 8, 4, 2, 1). \end{aligned}$$

From α , we see that for not quite one third of the time, Magnus finds himself in the state corresponding to the usual three ball “cascade” juggling pattern.

Remark 7. For other contexts in which the Stirling numbers of the second kind arise in conjunction with siteswap notation, see [5, 6]. A related class of numbers, the Eulerian numbers, also arise in the enumeration of siteswap patterns (see [4]).

3. PROOF OF THEOREM 5

The proof of Theorem 5 is not hard once we model the act of juggling randomly as a *Markov chain*: A discrete-time random process in which the transition probabilities depend only on the current state. To describe a Markov chain, we need to know the possible states Ω and the possible transitions among them. The latter we encode in a transition matrix P . We will denote this pair by $\text{MC}(\Omega, P)$. The (i, j) -th entry of the matrix P is to be interpreted as the probability of transitioning from the i -th state to the j -th state in one step.

In describing Magnus as taking a random walk on the states $St_{h,f}$ with throws chosen uniformly randomly, we have already described one Markov chain: $\text{MC}(St_{h,f}, P)$ where $P = (p_{\nu, \omega})_{\nu, \omega \in St_{h,f}}$ with

$$(5) \quad p_{\nu, \omega} = \begin{cases} 1, & \text{if } (\nu, \omega) \in \text{Edges}(G_{h,f}) \text{ and } \nu_1 = \smile, \\ \frac{1}{f+1}, & \text{if } (\nu, \omega) \in \text{Edges}(G_{h,f}) \text{ and } \nu_1 = \bullet, \\ 0, & \text{else.} \end{cases}$$

Suppose we have a Markov chain with states $\{\omega^1, \dots, \omega^r\}$ and transition matrix $P = (p_{i,j})$ such that the l -th power $P^l = (p_{i,j}^{(l)})$ of P satisfies $p_{i,j}^{(l)} > 0$ for $l \gg 0$. This condition on the matrix P means that, no matter how long Magnus has been juggling, he still has the opportunity to visit every state of $St_{h,f}$. $\text{MC}(St_{h,f}, P)$ satisfies this condition. A *probability vector* is a vector for which the sum of the entries is 1. The standard theorem of Markov chains we shall use is the following:

Theorem 8. ([7, 4.1.4]) *Given a Markov chain as above:*

1. *There exists a matrix A with $\lim_{l \rightarrow \infty} P^l = A$.*
2. *Each row of A is the same vector $\alpha = (\alpha_1, \dots, \alpha_r)$.*
3. *Each $\alpha_i > 0$.*
4. *The vector α is the unique probability vector satisfying $\alpha P = \alpha$.*
5. *For any probability vector π , $\pi P^l \rightarrow \alpha$ as $l \rightarrow \infty$.*

The number α_i can be interpreted as the fraction of time the Markov chain is expected to spend in the i -th state. Hence, Theorem 8 tells us how to answer our question: To find the fraction of time that Magnus is in state ω , we need only calculate the ω -th entry of the (normalized) left eigenvector for P with eigenvalue 1!

Example 9. We illustrate these ideas using the graph $G_{4,1}$ given in Figure 1. The transition matrix P simply consists of a $1/2$ for every edge except the unique edge leaving $\boxed{\smile \bullet \bullet \bullet}$, which must be taken. Notice that the sum of the entries in a row is 1. Below we list P and two of its powers.

$P =$	$\boxed{\bullet \bullet \bullet \smile}$	$\boxed{\bullet \bullet \smile \bullet}$	$\boxed{\bullet \smile \bullet \bullet}$	$\boxed{\smile \bullet \bullet \bullet}$
	1/2	1/2	1/2	1/2
	1/2			
	1/2			1/2
	1			
$P^2 =$	$\boxed{\bullet \bullet \bullet \smile}$	$\boxed{\bullet \bullet \smile \bullet}$	$\boxed{\bullet \smile \bullet \bullet}$	$\boxed{\smile \bullet \bullet \bullet}$
	1/2	1/4	1/4	
	1/2	1/4		1/4
	3/4	1/4		
	1/2	1/2		
$P^5 =$	$\boxed{\bullet \bullet \bullet \smile}$	$\boxed{\bullet \bullet \smile \bullet}$	$\boxed{\bullet \smile \bullet \bullet}$	$\boxed{\smile \bullet \bullet \bullet}$
	17/32	9/32	1/8	1/16
	17/32	1/4	5/32	1/16
	17/32	1/4	1/8	3/32
	9/16	1/4	1/8	1/16

Already for P^5 we see that the rows appear to be converging to the same vector. And, indeed, by Theorem 8 this vector is:

$$\begin{aligned}\alpha &= \frac{1}{\begin{Bmatrix} 5 \\ 2 \end{Bmatrix}} (\Delta(\boxed{\bullet\bullet\bullet\smile}), \Delta(\boxed{\bullet\bullet\smile\bullet}), \Delta(\boxed{\smile\bullet\bullet\bullet}), \Delta(\boxed{\smile\bullet\bullet\bullet})) \\ &= \frac{1}{15} (8, 4, 2, 1).\end{aligned}$$

Just over half of Magnus' time will be spent in the state $\boxed{\bullet\bullet\bullet\smile}$. In fact, this geometric distribution is what one would expect by examination of the graph $G_{4,1}$.

One can certainly prove Theorem 5 by guessing the vector α and then proving that it satisfies the conditions of Theorem 8. However, as is often the case in mathematics, keeping track of additional information enables us to write a simpler and more illuminating proof. So, we will construct an augmented Markov chain that is easy to analyze and from which we can obtain our original Markov chain by grouping states.

We will augment our states by specifying *which* balls are landing when. To this end, define a *Throwing/Landing-state* ("TL-state") to be an h -tuple $\widehat{\nu} = \boxed{\widehat{\nu}_1 \widehat{\nu}_2 \cdots \widehat{\nu}_h} \in \{\smile, 1, 2, \dots, h\}^h$ such that for each $1 \leq j, k, t \leq h$,

$$(6) \quad \widehat{\nu}_t \geq t \text{ or } \widehat{\nu}_t = \smile$$

and

$$(7) \quad \text{if } j \neq k, \text{ then } \widehat{\nu}_j - j \neq \widehat{\nu}_k - k.$$

The first condition ensures that any ball currently slated to land has, in fact, already been thrown. The second condition ensures that Magnus did not have to throw two balls at once in order to get into his current TL-state.

As is the case for states, we interpret $\widehat{\nu}_t = \smile$ to mean that no ball lands t seconds in the future. We interpret $\widehat{\nu}_t = i$ as meaning that the ball thrown $i - t$ seconds in the past will land t seconds in the future. (Alternatively, the ball landing t seconds in the future will have been in the air i seconds by the time it lands.) A TL-state completely captures the current status of Magnus. Let \widehat{St}_h denote the set of all TL-states of length h and $\widehat{St}_{h,f}$ the subset of TL-states with f \smile 's. The throwing operators Θ_j are extended in the obvious manner to act on TL-states.

Of course, each TL-state determines some landing state simply by forgetting how long each ball has been in the air. So, define a map $\pi_h : \widehat{St}_h \rightarrow St_h$ by $\pi_h(\widehat{\nu}) = \boxed{\nu_1 \nu_2 \cdots \nu_h}$ where

$$(8) \quad \nu_t = \begin{cases} \smile, & \text{if } \widehat{\nu}_t = \smile, \\ \bullet, & \text{if } \widehat{\nu}_t \in \{1, 2, \dots, h\}. \end{cases}$$

In addition to $\text{MC}(St_{h,f}, P)$, we now have another natural Markov chain to consider: $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$ where $\widehat{P} = (\widehat{p}_{\widehat{\nu}, \widehat{\omega}})_{\widehat{\nu}, \widehat{\omega} \in \widehat{St}_{h,f}}$ with

$$(9) \quad \widehat{p}_{\widehat{\nu}, \widehat{\omega}} = \begin{cases} 1, & \text{if } \widehat{\omega} = \Theta_0(\widehat{\nu}) \text{ and } \widehat{\nu}_1 = \smile, \\ \frac{1}{f+1}, & \text{if } \widehat{\omega} = \Theta_j(\widehat{\nu}) \text{ for some } j \text{ with } 1 \leq j \leq h, \widehat{\nu}_1 = \bullet \text{ and } \widehat{\nu}_{j+1} = \smile, \\ 0, & \text{else.} \end{cases}$$

Our approach to proving Theorem 5 will be to analyze this latter Markov chain. The first step will be to find the vector β of steady-state probabilities for $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$. The second step will be to show that "lumping" states of this chain recovers our original chain $\text{MC}(St_{h,f}, P)$. By analyzing this lumping process carefully, it will be clear how to obtain the result of Theorem 5 from our knowledge of the vector β .

Lemma 10. *The vector of all 1's (of length $|\widehat{St}_{h,f}|$) is a left eigenvector for \widehat{P} .*

Proof. To prove this, we need simply show that the transition matrix \widehat{P} is doubly stochastic (i.e., that the rows and columns sum to 1).

There are two cases to consider. In the first, either $\widehat{\nu}_t = \smile$ or $\widehat{\nu}_t > t$ for every t with $1 \leq t \leq h$. In this case, Magnus waited during the most recent beat. That means that the only precursor of $\widehat{\nu}$ is $\boxed{\smile \widehat{\nu}_1 \cdots \widehat{\nu}_{h-1}}$. On the other hand, if $\widehat{\nu}_t = t$ for some t , then he just made a height t throw. Since Magnus throws at most

one ball at any given time, there is at most one such t . By the discussion in the proof of Lemma 14 given below, the previous arc of this ball could have been any of $f + 1$ different heights. Each of these different heights leads to a different precursor. Since the out-degree of any state $\widehat{\nu}$ with $\widehat{\nu}_1 \neq \smile$ is $f + 1$, each of these edges has weight $\frac{1}{f+1}$. This proves that the columns of \widehat{P} also sum to 1. \square

In order to fully describe the vector β for $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$, we need to calculate $|\widehat{St}_{h,f}|$.

Lemma 11. $|\widehat{St}_{h,f}| = \binom{h+1}{f+1}$.

Proof. We proceed by constructing a bijection between elements of $\widehat{St}_{h,f}$ and partitions with $f + 1$ blocks of the set $[h + 1] := \{1, 2, \dots, h + 1\}$. An example of the correspondence we construct is illustrated in Figure 3.

We start by associating a graph $\Gamma_{\widehat{\nu}}$ with vertex set $[h + 1]$ to each $\widehat{\nu} \in \widehat{St}_{h,f}$. For $1 \leq t < i \leq h + 1$, we set $\{t, i\}$ to be an edge of $\Gamma_{\widehat{\nu}}$ if and only if there was a ball thrown $h + 1 - i$ seconds in the past that will be landing t seconds in the future. In terms of $\widehat{\nu}$, this means simply that $\widehat{\nu}_i = h + 1 + t - i$.

Let $\lambda(\widehat{\nu})$ be the set partition of $[h + 1]$ corresponding to the connected components of $\Gamma_{\widehat{\nu}}$. Our goal is to show that the map $\widehat{\nu} \mapsto \lambda(\widehat{\nu})$ is a bijection onto partitions of $[h + 1]$ with $f + 1$ blocks.

To begin, we examine the connected components of $\Gamma_{\widehat{\nu}}$. Suppose that $\{a, b\}, \{a, c\}$ are distinct edges of $\Gamma_{\widehat{\nu}}$. As they are both incident to the vertex a , all of a, b and c must be in the same connected component. It cannot be true that $a < b, c$ as then Magnus would be planning to catch two balls a seconds from now. Nor can it be true that $b, c < a$ as this would imply that a seconds ago Magnus threw two balls at once. We conclude that either $b < a < c$ or $c < a < b$. This implies that $\Gamma_{\widehat{\nu}}$ is a disjoint union of chains.

The number of edges in $\Gamma_{\widehat{\nu}}$ is $h - f$ since there is an edge for each j with $\widehat{\nu}_j \neq \smile$. So, the number of connected components in $\Gamma_{\widehat{\nu}}$ is the number of vertices $(h + 1)$ decreased by the number of edges $(h - f)$. This yields $f + 1$ connected components. Hence, $\lambda(\widehat{\nu})$ is a partition of $[h + 1]$ into $f + 1$ blocks as desired.

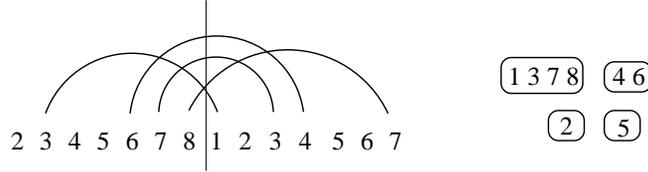


FIGURE 3. On the left we show the edges of $\Gamma_{\widehat{\nu}}$ for the TL-state $\widehat{\nu} = \boxed{6 \smile 46 \smile 7}$. On the right is the corresponding partition $\lambda(\widehat{\nu})$.

Now suppose that λ is a (set) partition of $[h + 1]$. We will associate a TL-state $\omega(\lambda) = \boxed{\omega(\lambda)_1 \cdots \omega(\lambda)_h}$ to λ . Consider a block $\{\alpha_1, \alpha_2, \dots, \alpha_m\}$ (indexed such that $\alpha_i < \alpha_j$ if $i < j$) of λ . If $m = 1$, then set $\omega(\lambda)_{\alpha_1} = \smile$. Otherwise, set $\omega(\lambda)_{\alpha_l} = h + 1 + \alpha_l - \alpha_{l+1}$ for each l with $1 \leq l < m$. Since $1 \leq \alpha_l < \alpha_{l+1} \leq h + 1$ for each l , $\omega(\lambda)$ is certainly an h -tuple in $\{\smile, 1, 2, \dots, h\}$. That (6) is satisfied follows from the inequalities $\alpha_l < \alpha_{l+1} \leq h + 1$. That $\omega(\lambda)$ satisfies (7) follows from the definition of $\omega(\lambda)_{\alpha_l}$ along with the fact that our blocks are disjoint. Hence, $\omega(\lambda)$ is, in fact, a TL-state. Furthermore, it can be checked that $\lambda \mapsto \omega(\lambda)$ is the inverse of the map $\widehat{\nu} \mapsto \lambda(\widehat{\nu})$. Therefore, every partition of $[h + 1]$ into $f + 1$ blocks is obtained as a $\lambda(\widehat{\nu})$ for a unique $\widehat{\nu} \in \widehat{St}_{h,f}$. This completes the proof. \square

Corollary 12. *In the process $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$, the fraction of time Magnus finds himself in state $\widehat{\nu}$ is $\frac{1}{\binom{h+1}{f+1}}$.*

We have shown that, when juggling randomly, Magnus is as likely to find himself in one TL-state as another. We will now consider a new random process running in parallel with $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$. For $\nu \in St_{h,f}$, define $[\nu] = \pi_h^{-1}(\nu) \in \widehat{St}_{h,f}$. When $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$ is in state $\widehat{\nu}$, we define our new process to be in the state $[\pi_h(\widehat{\nu})]$. It follows that its states are $\{[\nu]\}_{\nu \in St_{h,f}}$. This is a *lumped* process derived from $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$ by grouping together certain states.

Certainly Magnus wanders randomly among the states of this new process. We do not yet know, however, that this new process is a Markov chain. In particular, it is not clear that the probability of transitioning

from $[\nu]$ to $[\omega]$ is independent of how long Magnus has been in $[\nu]$. If it *were* dependent, then the transition probabilities would depend on previous states as well as the current state. Or, stated another way, the transition probability from $[\nu]$ to $[\omega]$ would depend on *which* state $\hat{\mu}$ of $[\nu]$ Magnus is currently in.

For each $\hat{\mu} \in [\nu]$, let $\hat{p}_{\hat{\mu},[\omega]}$ denote the probability of $\text{MC}(\widehat{St}_{h,f}, \widehat{P})$ transitioning from $\hat{\mu}$ to some element of $[\omega]$. Our above discussion suggests (see [7, 6.3.2] for a proof) that our new process will be a Markov chain if the probabilities $\hat{p}_{\hat{\mu},[\omega]}$ are equal for each $\hat{\mu} \in [\nu]$. Furthermore, the transition probabilities $r_{[\nu],[\omega]}$ of our new process will be these common values $\hat{p}_{\hat{\mu},[\omega]}$.

Lemma 13. *Fix $\nu, \omega \in St_{h,f}$. With the notation above, $\hat{p}_{\hat{\mu},[\omega]} = p_{\nu,\omega}$ for each $\hat{\mu} \in [\nu]$.*

Proof. If $p_{\nu,\omega} = 0$, then certainly all of the $\hat{p}_{\hat{\mu},[\omega]}$ are 0 too. Also, if $p_{\nu,\omega} = 1$, then Magnus must wait after any state $\hat{\mu} \in [\nu]$; hence $\hat{p}_{\hat{\mu},[\omega]} = 1$. So the only case to consider is when $p_{\nu,\omega} = \frac{1}{f+1}$. In this case, $\omega = \Theta_j(\nu)$ for some j with $1 \leq j \leq h$ and $\nu_1 = \bullet$. In particular, there is some j with $1 \leq j \leq h$ for which $\nu_{j+1} = \smile$ and $\omega_j = \bullet$. Suppose that Magnus transitions from $\hat{\mu}$ to a particular $\hat{\omega} \in [\omega]$. The previous equalities imply that $\hat{\mu}_{j+1} = \smile$ and $\hat{\omega}_j = j$. It follows that $\hat{\omega}$ is obtained from $\hat{\mu}$ by making a height j throw. This will be one of $f+1$ possible throws for Magnus as $\hat{\mu}_1 \neq \smile$; hence he will take this choice with probability $\frac{1}{f+1}$. As we have considered all cases, this completes the proof of the lemma. \square

Define a matrix $R = (r_{[\nu],[\omega]})_{\nu,\omega \in St_{h,f}}$ by setting $r_{[\nu],[\omega]} = p_{\nu,\omega}$. By the above lemma, we have obtained a new stochastic process $\text{MC}(\{[\nu]\}_{\nu \in St_{h,f}}, R)$ that can be identified in the obvious manner with $\text{MC}(St_{h,f}, P)$. It follows from Corollary 12 that the fraction of the time Magnus finds himself in a state ν with respect to the process $\text{MC}(St_{h,f}, P)$ is $\frac{|\nu|}{\sum_{\nu \in St_{h,f}} |\nu|}$. Hence, Theorem 5 follows from Lemma 14.

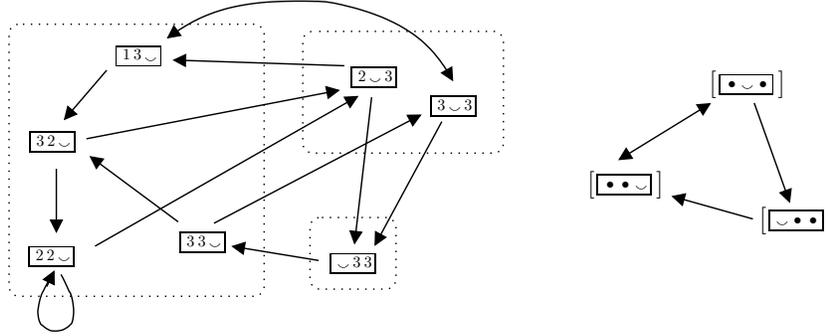


FIGURE 4. The Markov chain $\text{MC}(\widehat{St}_{3,1}, \widehat{P})$ along with the lumped process.

Lemma 14.

$$(10) \quad |[\nu]| = \Delta(\nu) = \prod_{t=1}^h (1 + \phi_t(\nu)).$$

Proof. Let $\hat{\nu} \in \widehat{St}_{h,f}$ and $\nu \in St_{h,f}$ such that $\pi_h(\hat{\nu}) = \nu$. Suppose $\nu_t = \bullet$ and that we know the value of $\hat{\nu}_i$ for each $i > t$. How do these values constrain the possibilities for $\hat{\nu}_t$? To start, note that the ball landing t seconds in the future was thrown at most $h-t$ seconds in the past. As Magnus must have already thrown this ball, there are at most $h-t+1$ possibilities for $\hat{\nu}_t$. In addition, any ball landing *more* than t seconds in the future must have been thrown *fewer* than $h-t$ seconds in the past as $\nu \in St_h$. Combining these two observations, we see that the number of possibilities for $\hat{\nu}_t$ is reduced precisely by the number of balls landing after it. In particular, the possible values for $\hat{\nu}_t$ depend only on the ν_i for $i > t$.

In conclusion, the number of possible values for $\hat{\nu}_t$ when $\nu_t = \bullet$ is

$$h-t+1 - |\{i > t : \nu_i = \bullet\}|,$$

which is just $(1 + \phi_t(\nu))$. Taking the product over all t yields the desired answer. \square

This completes the proof of Theorem 5.

Remark 15. Consideration of the powers of the transition matrix P might lead to new asymptotic formulae for the Stirling numbers of the second kind: Consider $G_{h,f}$ and the associated transition matrix P with the last row indexed by the state $\omega = \boxed{\smile \cdots \smile \bullet \cdots \bullet} \in St_{h,f}$. We have already seen that the rows of P^l converge to a probability vector α with $\alpha_\omega = 1/\{f+1\}^{h+1}$. Hence, for each state $\mu \in St_{h,f}$, we get a sequence

$$(p_{\mu,\omega}, p_{\mu,\omega}^{(2)}, \dots)$$

that converges to $1/\{f+1\}^{h+1}$.

4. OTHER MODELS

By now (unless you are a very fast reader), Magnus is getting tired. He is going to drop occasionally. Certainly any realistic juggling model should accommodate transitions to states with fewer balls. Of course, if we allow Magnus to juggle indefinitely, and he drops occasionally, he will eventually run out of balls. To keep things interesting, we will give his assistant, Sphagnum, the option of inserting a ball into Magnus' pattern whenever Magnus has a wait.

We will explore two such models in this section. In both models, Magnus can simply fail to catch a ball. When this happens, Magnus transitions from a state in $St_{h,f}$ to one in $St_{h,f+1}$. The difference between the two models lies in what is considered to be a legal throw. In the first model, Magnus (or his assistant) will only throw to heights that would not lead to two balls landing at the same time. In the second model, however, we allow this eventuality. Of course, as Magnus is well-known to have small hands, he cannot catch two at once, so he will have to ignore one of the balls. Such a throw will, in terms of transitions, be treated no differently than a drop.

These changes are easily incorporated into the framework we have developed. Now we define the state graphs corresponding to the aforementioned *add-drop model* (well known to college students) and *annihilation model*.

1. *Add-drop juggling* generalizes standard juggling by always allowing throws of height 0 and removing the restriction on $\nu_1 = \bullet$ for throws of positive height. (The throws of height 0 correspond to Magnus dropping or waiting; those of positive height when $\nu_1 = \smile$ correspond to Sphagnum helping.) So, our graph G_h^{a-d} has
 - (a) Vertices indexed by St_h .
 - (b) An edge from ν to ω whenever
 - (i) $\omega = \Theta_j(\nu)$ for some j , $1 \leq j \leq h$ such that $\nu_{j+1} = \smile$; or
 - (ii) $\omega = \Theta_0(\nu)$.

Figure 5 illustrates G_3^{a-d} .

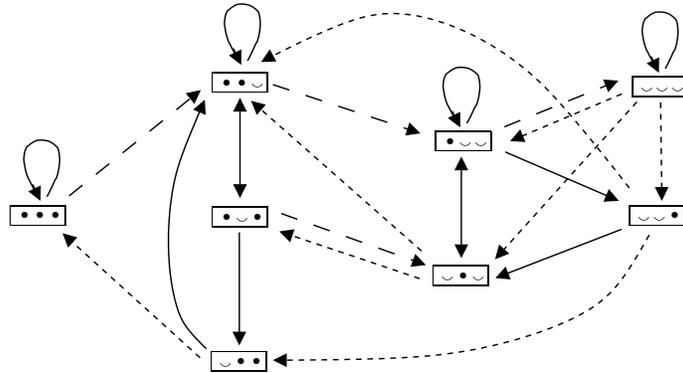


FIGURE 5. Solid arrows are regular throws; dashed are drops; dotted are insertions by Sphagnum.

2. In *Annihilation juggling*, we take the same state graph $G_h^{ann} = G_h^{a-d}$ as in the previous case. The difference will be in how we assign probabilities to the edges.

With our state graphs now defined, we can define the corresponding transition matrices. For $\nu \in St_{h,f}$, set $f'_\nu = |\{t : 2 \leq t \leq h \text{ and } \nu_t = \smile\}|$.

1. For G_h^{a-d} , set $Q = (q_{\nu,\omega})_{\nu,\omega \in St_h}$ with

$$q_{\nu,\omega} = \begin{cases} \frac{1}{f'_\nu+2}, & \text{if } (\nu, \omega) \in \text{Edges}(G_h^{a-d}), \\ 0, & \text{else.} \end{cases}$$

The denominator above comes from the fact that Magnus can either drop/wait or throw to any \smile occurring at ν_2 or later. So here, again, Magnus chooses uniformly from the available edges.

2. For G_h^{ann} , set $R = (r_{\nu,\omega})_{\nu,\omega \in St_h}$ with

$$r_{\nu,\omega} = \begin{cases} \frac{h-f'_\nu}{h+1}, & \text{if } (\nu, \omega) \in \text{Edges}(G_h^{ann}) \text{ and } \Theta_0(\nu) = \omega, \\ \frac{1}{h+1}, & \text{if } (\nu, \omega) \in \text{Edges}(G_h^{ann}) \text{ and } \Theta_0(\nu) \neq \omega, \\ 0, & \text{else.} \end{cases}$$

Here, you should envision Magnus (possibly with the help of Sphagnum) throwing to *any* height ($\leq h$) at random, regardless of whether or not a ball has just landed. If there is already a ball landing at the time he is throwing to, the latter ball gets annihilated.

In order to describe the steady-state probabilities for the add-drop model, we must introduce the Bell numbers. The h -th Bell number, B_h , is defined by $B_h = \sum_{i=0}^h \{h \atop i\}$. For these two models, the analogue of Theorem 5 is

Theorem 16. *Let $\nu \in St_{h,f}$.*

1. *In the add-drop model, Magnus spends*

$$\frac{\Delta(\nu)}{B_{h+1}}$$

of his time in state ν .

2. *In the annihilation model, Magnus spends*

$$\frac{\frac{h!}{(h-f)!} \Delta(\nu)}{(h+1)^h}$$

of his time in state ν .

As the denominators for our probabilities in the standard and add-drop model have natural combinatorial interpretations, the reader may be wondering if such an interpretation exists for the $(h+1)^h$ occurring in the annihilation model. Indeed, $(h+1)^h$ counts the number of labeled rooted trees with h nodes (see [10, 5.3.2] for details).

Example 17. Below we give the transition matrix for G_3^{a-d} :

	1/2	1/2						
		1/3	1/3		1/3			
			1/3	1/3		1/3		
	1/2	1/2						
					1/4	1/4	1/4	1/4
			1/3	1/3	1/3			
						1/3	1/3	
					1/4	1/4	1/4	1/4

From Theorem 16.1, we obtain the probability vector

$$\alpha = \frac{1}{15} (1, 4, 2, 1, 3, 2, 1, 1).$$

This gives the fraction of the time Magnus spends in each of the eight possible states (the ordering of α is the same as that of the transition matrix).

As the steps in the proof of Theorem 16 are analogous to those found in the proof of Theorem 5, we do not detail them here. As a consolation to the reader, though, we mention two other families of juggling models which, although perhaps less physically realistic than the ones we have described, do lead to interesting answers to our Random Question.

1. *Multiplex juggling*: There is no real reason (assuming you find a juggler with bigger hands) to disallow the catching and throwing of more than one ball at a time. The standard and add-drop models can both be reinterpreted with this condition relaxed.
2. *Infinite juggling*: Again, given a juggler in better shape than Magnus, there is no reason to place a maximum on legal throw heights. In the standard juggling model, placing a probability of $(p-1)/p^j$ for the throw to the j -th next available height yields tractable calculations for any $p > 1$.

Rather than choosing geometric weights for the throws, one could choose probabilities from any convergent infinite series. For example, why not choose $6/(\pi^2 j^2)$ for the j -th available throw height? Or even $90/(\pi^4 j^4)$? The sky's the limit.

5. ACKNOWLEDGEMENTS

I would like to thank Michael Schneider for helpful advice on Markov chains and Joe Buhler for help in sorting out who should be credited with the idea of siteswap notation.

REFERENCES

- [1] <http://www.jugglingdb.com/>.
- [2] http://directory.google.com/Top/Arts/Performing_Arts/Circus/Juggling/Software/.
- [3] <http://www.juggling.org/programs/java/MAGNUS/>.
- [4] Joe Buhler, David Eisenbud, Ron Graham, and Colin Wright. Juggling Drops and Descents. *Amer. Math. Monthly*, 101:507–519, 1994.
- [5] R. Ehrenborg and M. Readdy. Juggling and applications to q -analogues. In *Proceedings of the 6th Conference on Formal Power Series and Algebraic Combinatorics (New Brunswick, NJ, 1994)*, volume 157, pages 107–125, 1996.
- [6] L. Kamstra. Juggling polynomials. Technical report, Centrum voor Wiskunde en Informatica, 2001.
- [7] J. Kemeny and J. Snell. *Finite Markov chains*. Van Nostrand, Princeton, N.J., 1960.
- [8] Burkard Polster. *The Mathematics of Juggling*. Springer Verlag, 2002.
- [9] R. Stanley. *Enumerative Combinatorics*, volume 1. Wadsworth & Brooks/Cole, Belmont, CA, 1986.
- [10] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*. Cambridge University Press, Cambridge, 1999.