ON THE NUMBER OF FAILURES UNTIL A FIRST SUCCESS IN n BERNOULLI TRIALS CONTAINING m = 1, ..., n SUCCESSES

William E. Wright, Dept. of Computer Science
Sakthivel Jeyaratnam, Dept. of Mathematics
Southern Illinois University Carbondale
Carbondale, IL 62901

Abstract. We describe a random variable $D_{n,m}$, $n \geq m \geq 1$, as the number of failures until the first success in a sequence of n Bernoulli trials containing exactly m successes, for which all possible sequences containing m successes and n-m failures are equally likely. We give the probability density function, the expectation, and the variance of $D_{n,m}$. We define a random variable D_n, $n \geq 1$, to be the mean of $D_{n,1}, \ldots, D_{n,n}$. We show that $E[D_n]$ is a monotonically increasing function of n and is bounded by ln n. We apply these results to a practical application involving a video-on-demand system with interleaved movie files and a delayed start protocol for keeping a balanced workload.

Key words. Bernoulli trials, video on demand, movie on demand, interleaved files

1. A Practical Problem

In a video-on-demand (VOD) system, customers are able to watch a movie on their television at home, using a VOD terminal instead of a VCR. Large files of compressed digitized movies are stored in an array of magnetic disks on the VOD server. The server transmits blocks of the movie file over a communication line to the terminal, on a just-in-time basis. At any time, the system may be serving many users (e.g., 3,000). Each user is treated as a movie stream. Different streams may use the same movie file, with different start times. Wright [7] describes an efficient model for the system.

The server reads the disks in rounds, reading the necessary next blocks for all of the movie streams. A block might typically generate one or two seconds of playing time, and the rounds need to take less time than that in order to keep up with the "consumption" of the digitized data at the terminals. Of all the blocks (3,000, say) that need to be read during a round, each disk must read whichever ones are stored on it.

An efficient pattern for storing the movies on the disks uses interleaving, whereby each movie file has its blocks spread in a regular sequence over all the disks. Figure 1 illustrates interleaving in which there are 100 disks and 500 available movies. The figure shows (for simplicity) each movie having its first block on disk 1, but the starting disks for the movies would probably be randomized.

The purpose of interleaving is to balance the workload evenly over all the disks. Each movie stream accesses all of the disks in sequence, hence all of the disks are used evenly, on the average. If each movie file were stored instead on a single disk, then the disks storing the most popular movies would be overworked and disks storing unpopular movies would be underutilized. (Balancing popular with unpopular movies on the same disk would help, but it would generally not be nearly sufficient to even out the workload. It would also be vulnerable to unreliable and changing predictions of movie popularity.)

	Disk 1	Disk2	...	Disk 100
Movie 1	Block 1	Block 2	...	Block 100
	Block 101	Block 102	...	Block 200
	...			
...				
Movie 500	Block 1	Block 2	...	Block 100
	Block 101	Block 102	...	Block 200
	...			

Figure 1.
Interleaved Files

Although interleaving balances the workload evenly on the average, it does not guarantee a balance during each round. For example, if there are 2,864 movie streams and 100 disks, then the average number of block reads per round per disk is 28.64. However, due to the randomness of when each stream started (and which disk it started on), some disks might have to read 34 blocks during a round and others only 25, say. The best possible balance would occur if exactly 64 disks were serving 29 streams during a round, and 36 disks were serving 28.

In general, the best balance is attained when

$$max = \lceil s/n \rceil \quad \text{and} \quad min = \lfloor s/n \rfloor \qquad (1)$$

where s is the number of streams, n is the number of disks, and max (min) is the maximum (minimum) number of streams served by all of the disks. In this case, the number of disks serving max and min streams during a round is given by

$$n_{max} = s - min \cdot n$$
$$n_{min} = n - n_{max}$$

Note that we must have

$$max - min = 0 \quad \text{if s is a multiple of n}$$
$$max - min = 1 \quad \text{otherwise}$$

Balance is achieved when max and min differ by at most 1.

An important goal of a VOD system design is to make the system capable of serving as many simultaneous streams as possible without a "starvation," i.e., a terminal whose buffer empties out before the next block arrives. Since each disk has time to read only a limited number of blocks during a round, it follows that the design should try to keep max to a minimum, which is accomplished by ensuring that (1) holds.

Fortunately there is a simple strategy for enforcing this constraint in the presence of new movie requests. The strategy is to delay the start of a new stream, as necessary, until the disk on which it starts will be serving the minimum number of streams. Using the example above, suppose the system is prepared to start a new movie stream on the current round. If the starting disk for the stream has 28 other streams assigned to it for the round, then it can start the new stream, giving it 29. However, if it already has 29 streams, then the system delays the start until the next round and repeats the test.

Using interleaving, the streams for one disk on one round are passed on to the next disk for the next round. Accordingly, the policy is to continue delaying the start of the new stream until the first round for which the starting disk is assigned only 28 other streams. This process guarantees that each new stream will start when its disk is serving only 28 other streams, until the total number of streams reaches 2900. At this time min becomes 29, max becomes 30 (after one more new stream), and the process repeats.

This straight-forward strategy can maintain a balanced workload over all the disks during every round, regardless of movie popularity. However, it does raise (at least) one question: how many rounds will a new user have to wait, on average, for his movie to start? That number is a random variable, and we wish to analyze it in this paper.

2. Mathematical Background

Several discrete probability distributions have been defined based on repeated Bernoulli trials. Some examples presented in Parzen [6] are as follows:

1. The binomial distribution with parameters p and n can be characterized as the number of successes in $n \geq 1$ independent Bernoulli trials in which the probability of success in each trial is p, $0 \leq p \leq 1$.
2. The geometric distribution with parameter p can be characterized as the number of trials until the first success in a sequence of independent Bernoulli trials for which the probability of success in each trial is p.
3. The negative binomial distribution with parameters p and m can be characterized as the number of failures until the m-th success ($m \geq 1$) in a sequence of independent Bernoulli trials for which the probability of success in each trial is p.

We analyze here a probability distribution that can be characterized in terms of a sequence of Bernoulli trials which are not independent. For convenience of reference in this paper, we need to name the distribution, and have rather arbitrarily chosen to refer to it as the "delta distribution." Accordingly, let n and m be integers with $1 \leq m \leq n$. A random variable is said to have the delta distribution with parameters n and m if it can be characterized as the number of failures until the first success in a sequence of n Bernoulli trials containing exactly m successes, for which all possible sequences containing m successes and n-m failures are equally likely. The sequence of trials is clearly not independent since, for example, the outcome of the last trial can be determined exactly from the outcomes of the first n-1 trials.

It is interesting to note that the delta distribution has a combination of characteristics from the binomial, geometric, and negative binomial distributions. Like the binomial distribution, it has a fixed number of trials. Like the geometric distribution, the random variable is the number of trials until the first success. Like the negative binomial distribution, there are m successes in the sequence of trials. Unlike these three distributions, the delta distribution does not have an explicit parameter p, although it is obvious that the probability of success in each trial is $p = m/n$.

It is also possible to characterize the delta distribution in other ways. Feller [1] describes a class of distributions based on sampling balls from an urn. He assumes that an urn originally contains b black balls and r red balls. A ball is drawn at random and replaced. Moreover, c balls of the color drawn and d balls of the opposite color are added to the urn, for some parameters c and d, and the procedure is repeated. If $c = d = 0$, we get independent and identical Bernoulli trials. If we look for runs of the form bb...br, we get the geometric distribution. For runs of the same form with $c = -1$ and $d = 0$, we get the proposed delta distribution.

The distribution of runs of a given length was studied by Mood [5]. Statistics relating to runs of both identical and non-identical independent Bernoulli trials were presented by Fu and Koutras [2]. The number of 0 bits until the first 1 bit in a random bit string containing exactly m 1 bits and exactly n-m 0 bits, has the delta distribution. In this context, a 0 bit is interpreted as a failure (in a Bernoulli trial) and a 1 bit as a success. For a treatment based on strings, see Guibas and Odlyzko [3].

3. Analysis of the Distribution

Let $D_{n,m}$ be a delta random variable with parameters $n \geq m \geq 1$. We use C(a,b) to denote the number of combinations of a objects taken b at a time. For d an integer with $0 \leq d \leq n-m$,

$P[D_{n,m} = d]$ = P[the first d trials are failures and the (d+1)-th trial is a success]

$$= \frac{C(n-d-1,m-1)}{C(n,m)}$$

Therefore the probability mass function $f_D(d)$ is given by

$$f_D(d) = \frac{C(n-d-1,m-1)}{C(n,m)} \qquad \text{if } 0 \le d \le n\text{-}m \qquad (2)$$
$$= 0 \qquad \text{otherwise}$$

To show that (2) defines a valid density function, we prove the following:

Theorem 1. $\displaystyle\sum_{d=0}^{n-m} \frac{C(n-d-1,m-1)}{C(n,m)} = 1$

Proof: (by induction on n)

If n = m, the result is clearly true. Assume it is true for n = k ≥ m. Then

$$\sum_{d=0}^{k+1-m} \frac{C(k+1-d-1,m-1)}{C(k+1,m)} = \frac{C(k,m-1)}{C(k+1,m)} + \sum_{d=1}^{k+1-m} \frac{C(k+1-d-1,m-1)}{C(k+1,m)}$$
$$= \frac{m}{k+1} + \sum_{d'=0}^{k-m} \frac{C(k-d'-1,m-1)}{C(k+1,m)}$$
$$= \frac{m}{k+1} + \frac{C(k,m)}{C(k+1,m)} = \frac{m}{k+1} + \frac{k+1-m}{k+1} = 1$$

For fixed n > m ≥ 1 and d = 0, 1, …, n-m-1,

$$f_D(d) - f_D(d+1) = \frac{(n-d-2)!(m-1)}{(m-1)!(n-m-d)!C(n,m)}$$

Therefore, as a function of d, $f_D(d)$ is decreasing if m > 1 and constant if m = 1. Figure 2 is a graph of $f_D(d)$ for n = 10 and m = 3.

We now wish to determine the expectation and variance of $D_{n,m}$.

Lemma: For any positive integer j, $E[D^j_{n+1,m}] = \frac{n-m+1}{n+1} E[(D_{n,m} + 1)^j]$

Proof: (by induction on n)

If n = m, the result is clearly true, since $D_{n,n} \equiv 0$. Assume it is true for n = k ≥ m. Then

$$E[D^j_{k+1,m}] = E[D^j_{k+1,m} \mid \text{first trial is a failure}] \cdot P[\text{first trial is a failure}]$$
$$+ E[D^j_{k+1,m} \mid \text{first trial is a success}] \cdot P[\text{first trial is a success}]$$
$$= E[(D_{k,m}+1)^j] \frac{k-m+1}{k+1} + 0 \cdot \frac{m}{k}$$
$$= \frac{k-m+1}{k+1} E[(D_{k,m} + 1)^j]$$

Thus the formula holds for n = k+1, and by induction it holds for all n ≥ 1.

Theorem 2. $E[D_{n,m}] = \frac{n-m}{m+1}$

Proof: (by induction on n)

If n = m, the result is clearly true, since $E[D_{n,n}] = 0$. Assume it is true for n = k ≥ m. Then by the lemma

$$E[D_{k+1,m}] = \frac{k-m+1}{k+1} E[D_{k,m} + 1]$$

$$= \frac{k-m+1}{k+1}\left(\frac{k-m}{m+1}+1\right)$$
$$= \frac{k+1-m}{m+1}$$

Thus the formula holds for n = k+1, and by induction it holds for all n ≥ 1.

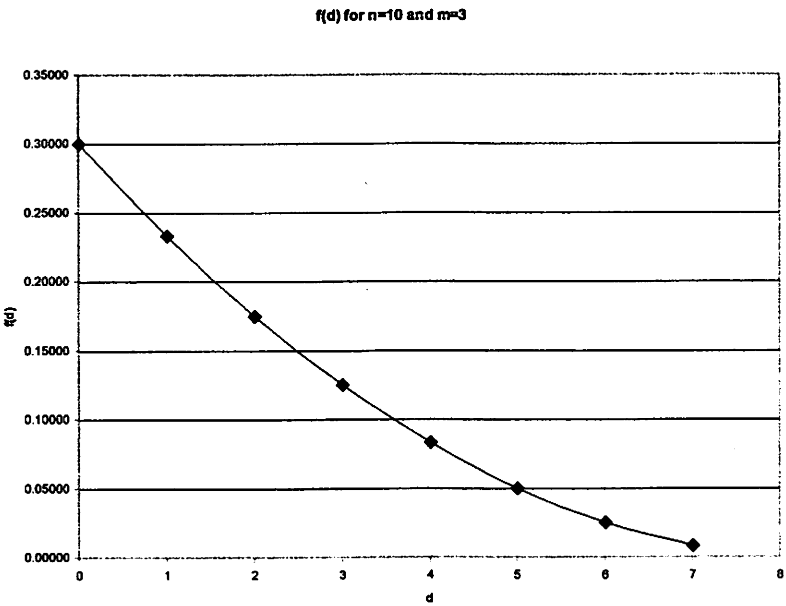

Figure 2.
Probability Mass Function of $D_{10,3}$

Theorem 3. $E[D^2_{n,m}] = \frac{(n-m)(2n-m)}{(m+1)(m+2)}$

Proof: (by induction on n)

If n = m, the result is clearly true, since $E[D^2_{n,m}] = 0$. Assume it is true for n = k ≥ m. Then by the lemma

$$E[D^2_{k+1,m}] = \frac{k-m+1}{k+1} E[(D_{k,m} + 1)^2]$$
$$= \frac{k-m+1}{k+1}\left(\frac{(k-m)(2k-m)}{(m+1)(m+2)} + \frac{2(k-m)}{m+1} + 1\right)$$
$$= \frac{(k+1-m)(2k+2-m)}{(m+1)(m+2)}$$

Thus the formula holds for n = k+1, and by induction it holds for all n ≥ 1.

From Theorems 2 and 3 we get

$$\mathrm{Var}[D_{n,m}] = \frac{(n-m)(2n-m)}{(m+1)(m+2)} - \left(\frac{n-m}{m+1}\right)^2 = \frac{m(n+1)(n-m)}{(m+1)^2(m+2)}$$

4. The Mean of $D_{n,1}, \ldots, D_{n,n}$

For $n \geq 1$ we define the random variable D_n to be the mean of $D_{n,1}, \ldots, D_{n,n}$, i.e.,

$$D_n = \frac{1}{n}\sum_{m=1}^{n} D_{n,m}$$

Then $$E[D_n] = \frac{1}{n}\sum_{m=1}^{n} \frac{n-m}{m+1}$$

Table 1 gives the value of $E[D_n]$ for several sample values of n.

Table 1.
Sample Values of $E[D_n]$

n	1	2	5	10	50	100	200	500	1000
$E[D_n]$	0	.25	.74	1.2219	2.5892	3.2393	3.9074	4.8064	5.4930

Theorem 4. $e(n) = E[D_n]$ is a monotonically increasing function of n.
Proof:

We need only show that $e(n+1)-e(n) > 0$ for $n \geq 1$. We have

$$\begin{aligned} e(n+1)-e(n) &= \frac{1}{n+1}\sum_{m=1}^{n+1} \frac{n+1-m}{m+1} - \frac{1}{n}\sum_{m=1}^{n} \frac{n-m}{m+1} \\ &= \frac{1}{n+1}\sum_{m=1}^{n} \frac{n+1-m}{m+1} - \frac{1}{n}\sum_{m=1}^{n} \frac{n-m}{m+1} \\ &= \sum_{m=1}^{n}\left(\frac{n+1-m}{(m+1)(n+1)} - \frac{n-m}{(m+1)n}\right) \\ &= \sum_{m=1}^{n} \frac{m}{(m+1)(n+1)n} \\ &> 0 \end{aligned}$$

Theorem 4 enables us to conclude, for example, that not only is $E[D_{1000}] = 5.4930$ but $E[D_n] < 5.4930$ for every $n < 1000$.

Theorem 5. $E[D_n] \le \ln n$ for $n \ge 1$, i.e., $\ln n$ is a bounding function for $E[D_n] = e(n)$.

Proof:

$$e(n) = \frac{1}{n}\sum_{m=1}^{n}\frac{n-m}{m+1}$$

$$\le \frac{1}{n}\sum_{m=1}^{n}\frac{n}{m+1}$$

$$= H_n - 1, \quad \text{where } H_n \text{ is the harmonic series } \tfrac{1}{1}+\tfrac{1}{2}+\ldots+\tfrac{1}{n}.$$

$$H_n < \ln n + .577 + \tfrac{1}{2n} \qquad \text{from Knuth [4]}$$

$$e(n) \le \ln n + .577 + .25 - 1 < \ln n \qquad \text{for } n \ge 2$$

$$e(1) = 0 = \ln 1,$$

$$e(n) \le \ln n \qquad \text{for } n \ge 1$$

Table 2 gives some sample values of ln n.

Table 2.
Sample Values of ln n

n	100	500	1000	5000	20000
ln n	4.605	6.215	6.908	8.517	9.903

Thus, for example, the expectation of D_n is less than 10, even if n is as large as 20,000.

5. Application to the Video-on-Demand Problem

Referring back to the example and notation in Section 1, let $k = s - \min$ and $m = n - k$. Then k is the number of disks serving the maximum (e.g., 29) number of streams during a round, and m is the number serving the minimum (e.g., 28). Define a Bernoulli trial to be a success if the starting disk for a new movie request is serving min streams during the current round, and a failure otherwise. The number $R_{n,m}$ of rounds the VOD system must wait, under the delayed-start policy, is then the number of failures until the first success. The m disks serving the minimum are distributed randomly among the n disks, with all possible combinations equally likely. Hence, $R_{n,m}$ has a delta distribution with parameters n and m and

$$E[R_{n,m}] = \frac{n-m}{m+1}$$

$$Var[R_{n,m}] = \frac{m(n+1)(n-m)}{(m+1)^2(m+2)}$$

from Theorems 2 and 3.

In realistic VOD systems, it is reasonable to assume that, over most time periods, k and m take on the values 0, 1, ..., n-1 with equal likelihood. Hence the average value R_n of $R_{n,m}$ over time is given by

$$R_n = \frac{1}{n}\sum_{m=1}^{n} R_{n,m}$$

and the results of Theorems 4 and 5 and Tables 1 and 2 apply to R_n. For example, $E[R_n] \leq \ln n$ for all n, $E[R_n] < 5.4931$ for $n \leq 1{,}000$, and $E[R_n] < 10$ for $n \leq 20{,}000$.

References

1. Feller, W. *An Introduction to Probability Theory and Its Applications*, John Wiley, New York, 1968, pp. 119-120.
2. Fu, J.C. and Koutras, M.V. Distribution theory of runs: a Markov chain approach. **J. Amer. Statis. Assoc. 89**, 1994, pp. 1050-1058.
3. Guibas, L.J. and Odlyzko, A.M. String overlaps, pattern matching, and nontransitive games. **J. Combin. Theory A30**, 1981, pp. 183-208.
4. Knuth, D.E. *The Art of Computer Programming, Vol. 1, Fundamental Algorithms*. Addison-Wesley, Reading, MA, 1969, p. 74.
5. Mood, A.M. The distribution theory of runs. **Ann. Math. Statis. 11**, 1940, pp. 367-392.
6. Parzen, E. *Stochastic Processes*. Holden-Day, San Francisco, 1962, pp. 13-15.
7. Wright, W.E. An efficient video-on-demand model. **Computer 34**, 5, 2001, pp. 64-70.