Probabilistic packet transmission through a limited-lifetime deletion channel with arbitrary deletion probability

V. Hunter Adams

Abstract—This paper considers the transmission of a packet through a binary deletion channel with an unknown and arbitrarily high deletion probability. The deletion positions are unknown to both sender and receiver, and the number of bits received before channel closure is unknown to both sender and receiver. The number of received bits is in the range 1-100. This paper argues that, for such a channel, an encoding method that enables the calculation of probability masses for all possible packets is preferred to one which attempts perfect reconstruction of the transmitted packet. It then presents and analyzes such an encoding method. The presented method encodes information on the probability of receiving a 1 vs. a 0, rather than in the 1's and 0's themselves. The receiver decodes the information using Bayesian inference. Rather than transmitting the packet of interest, the encoding method transmits information about the packet of interest in order to guarantee information transfer for any nonzero number of bits received. This paper analyzes the statistics of convergence of this estimator for all numbers of received bits from 0 to 100.

Index Terms—deletion channel, probabilistic transmission, spacecraft, starshot, estimation, encoding.

I. INTRODUCTION

THIS paper presents and analyzes a method for probabilistically encoding a packet for transmission through a binary deletion channel with an arbitrarily high deletion probability and unknown channel lifetime. The deletion channel takes an input sequence of N bits, X_i. The value of N (the number of bits that the sender transmits before the channel closes) is unknown to both the sender and receiver and is assumed to be small (< 100 bits). Each bit can be deleted with a probability p, which is unknown and may be arbitrarily close to 1. Neither the sender nor the receiver know the positions of the deleted bits. The output sequence, Y_i, is the sequence of the X_i which were not deleted, in the correct order and with no errors. The length of Y_i is assumed to be at least 1 bit and less or equal to N bits with an expected length of $N \cdot (1-p)$. The transmission is not decoded to a single packet, but instead to a probability density function that assigns a probability density to every possible packet.

Neither the sender nor the receiver know when the channel will close or which bits have been deleted. For that reason, each bit should include information about the entire packet that the sender is attempting to send through the channel. Otherwise there will be some minimum number of bits required for the receiver to recover any useful information about the packet.

Lecturer of Mechanical and Aerospace Engineering, Cornell University

This paper presents a method for encoding information about the entire packet in each bit. Using the proposed encoding, any number of bits recovered by the receiver (no matter how few) improves the receiver's knowledge of the entire transmitted packet.

This paper proposes to encode information in the following way. The sender and receiver both assume a fixed, known packet length. The packet to be transmitted is interpreted as an integer. The ratio between this integer and the largest integer storable in the packet length is interpreted as a probability, x. The sender uses x to decide whether to transmit a binary 1 or 0. Each received bit is modeled as an independent and identically distributed Bernoulli random trial. The bit is 1 with probability x and 0 with probability $(1 - x)$. The probability density function for x is calculated by the receiver via Bayesian inference.

By encoding information in this way, every bit adds to the receiver's knowledge of the transmitted information. Each subsequent bit improves the estimate for x, thereby tightening the probability density function for possible transmissions around fewer packets. The statistics of convergence are discussed at length in this paper. Thus, even for a channel which may prevent reconstruction of the particular packet which the sender transmitted, the proposed encoding scheme allows the receiver to assign probability masses to each of every possible packet.

The scope of application for this encoding method is limited to channels with the properties enumerated below. These are deletion channels with properties that do not guarantee the intact receipt of any packet larger than one bit, and that do not guarantee the receipt of enough bits to perfectly reconstruct packets longer than 1 bit from partial packets. For such channels, perfect reconstruction of a transmitted packet is sacrificed in lieu of probability masses over all possible packets in order to guarantee information transfer through the channel.

Channel properties
   1) The channel is a binary deletion channel.
   2) The channel has unknown and arbitrary high deletion probability, p.
   3) Deletion positions unkown to both sender and receiver.
   4) The number of bits received before channel closure, n, is unknown to both sender and receiver and assumed greater than or equal to one bit and less than or equal to 100 bits.
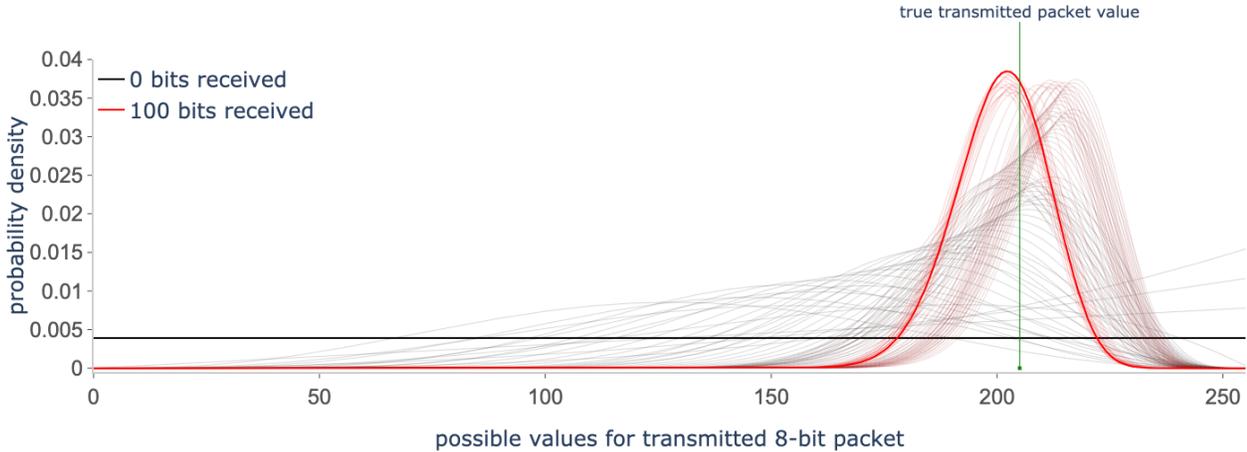
Fig. 1. Probability density function for the value of the transmitted packet based on a number of received bits that varies from 0 (black) to 100 (red). Figure represents a particular instantiation of a random process.

II. MOTIVATION

The channel described in the introduction models the channel which can likely be expected for the Breakthrough Starshot Project, which proposes to send a number gram-scale spacecraft to the star Proxima Centauri — 4.22 light years from Earth. The spacecraft will use lasers to transmit information to the Earth. Photons from the lasers will be collected in telescopes on Earth and decoded into a bit string [1]. At these distances, tremendous amounts of data will be required in order to resolve even a single bit from background noise, and many bits will be deleted from the channel. Every recovered bit has far greater value than in conventional communications channels. The channel to each of the spacecraft is fragile and could be severed at any time, and many of the bits that the spacecraft send will not be recoverable by the receiver. The encoding method proposed in this paper is intended to maximize the scientific impact of every received bit.

This paper proposes that the spacecraft used for the Breakthrough Starshot Project should not attempt to transmit a packet of information as a string of ones and zeros. If only a handful of bits are recovered from the deletion channel, there may be no recoverable information about the original packet. Instead, this paper proposes to encode the packet information in the probability of transmitting a 1 vs. a 0, as described in Section III. In doing so, any bit recovered from the channel adds information about the original transmitted packet. The packet is recovered with increasing certainty (the precise statistics of which are discussed in a Section VI) as more bits are recovered from the channel. The objective should not be to maximize data rate, but instead to guarantee information transfer.

III. RELATED WORK

Binary deletion codes were introduced as a concept for study by Levenshtein in 1966 [2]. Since then, much of the research associated with deletion channels has been toward bounding the channel capacity, the precise calculation of which

remains an unsolved problem. Diggavi et al present the first non-trivial upper bound for the capacity of a probabilistic deletion channel in [3]. In [4], Fertonani and Duman use the upper bound from [3] as a benchmark, and calculate a new state-of-the-art for the upper-bound channel capacity for most deletion probabilities. To my knowledge, theirs is the lowest bound for the upper limit of the channel capacity. Similar efforts have been made in calculating the lower bound, with the current state-of-the-art calculated by Drinea and Mitzenmacher in [5]. This paper does not attempt to bound the capacity of the deletion channel of interest, but instead presents an encoding strategy for guaranteeing communication of information through that channel. There has been related research on encoding strategies for deletion channels.

In [6], Kanoria and Montanari develop an optimal code for a binary deletion channel with a small deletion probability. They do so by beginning with a channel with a deletion probability of zero, and assuming that a channel with a small deletion probability has a capacity that varies smoothly with the deletion probability. They compute the capacity as a series expansion over small deletion probability and find an input distribution that achieves capacity up the the first three terms in this expansion. The work in this paper departs from theirs in that the deletion probability may be arbitrarily large. It departs from the body of work surrounding deletion channels in other ways too.

This work departs from existing work in that the objective is not to maximize data rate through a permanent channel. Instead, the objective to guarantee the communication of information *about* a packet of interest through a channel with an unknown and arbitrarily high deletion probability, unknown lifetime, and through which an unknown and small (1-100) number of bits can be recovered. For such a channel, intact receipt of the packet of interest cannot be guaranteed before the channel closes. Furthermore, a sufficient number of partial packets for reconstructing the packet of interest cannot be guaranteed before the channel closes. This calls for an encoding method that enables the assignment of probability

masses to all possible packets rather than one which attempts perfect reconstruction of the packet.

IV. ENCODING

The sender interprets the packet to be transmitted as an integer. The length of the packet is fixed and known by both the sender and receiver. The sender finds the ratio, x, between the packet integer value, v, and the maximum integer value that can be contained in a number of bits equal to the packet length, M. M is 256 for 8-bit packets, for example. This calculation is shown in eqn. 1.

$$x = \frac{v}{M} \tag{1}$$

The sender interprets this ratio, which will be in the range [0,1], as a probability. Before transmitting each bit, a uniform random number generator onboard the sender generates a number in the range [0,1]. If the number is less than x, then the sender transmits a 1. If the number is greater than x, the sender transmits a 0. The receiver gathers these bits and uses Bayesian inference to assign a probability density to each of every possible packet value, v.

V. DECODING

The transmission is not decoded to a single packet, but instead to a probability density function that assigns a probability density to every possible packet value. This is accomplished through a straightforward application of Bayes' Rule, shown in eqn. 2. In eqn. 2, $p(k|x)$ is the probability of receiving k ones conditioned on the probability of sending a 1. $p(x)$ is the prior distribution on x, the probability of receiving a 1 (which may or may not be uniform, depending on the application). $p(k)$ is the marginal probability of receiving k 1's.

$$p(x|k) = \frac{p(k|x)p(x)}{p(k)} \tag{2}$$

The receiver sporadically receives bits (1 or 0). Each bit has probability x of being 1 and probability $(1-x)$ of being 0. Thus, each received bit is a Bernoulli random trial. A sequence of such bits is a Bernoulli process with the probability mass function given by the binomial distribution shown in eqn. 3. The binomial distribution, with parameters, n and k, is the discrete probability distribution of the number of successes (1's) in a sequence of n experiments (received bits). This distribution, in other words, yields the probability of k 1's received in n total bits, conditioned on the probability of receiving a 1. This is the first term in the numerator of Bayes' Rule (eqn. 2).

$$p(k|x) = \frac{n!}{k!(n-k)!}x^k(1-x)^{n-k} \tag{3}$$

One might choose any one of a number of distributions to represent the prior distribution on x, $p(x)$. However, the beta distribution is the conjugate prior for the binomial distribution. Thus, by choosing to describe $p(x)$ with a beta distribution, we know that $p(x|k)$, the posterior distribution, will also be a beta function. α and β in eqn. 4 are hyper parameters that

affect the shape of the distribution. When $\alpha = \beta = 1$, the beta distribution reduces to a uniform distribution (all packets equally likely). Other choices of α and β place more weight on various values for x.

$$p(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}x^{\alpha-1}(1-x)^{\beta-1} \tag{4}$$

The marginal probability for receiving k 1's is calculable using the chain rule and the known solution to an Euler integral, as shown in eqn. 5.

$$p(k) = \int_k p(k|x)p(x)dx$$
$$= \frac{n!}{k!(n-k)!} \cdot \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \cdot \frac{\Gamma(k+\alpha)\Gamma(n-k+\beta)}{\Gamma(n+\alpha+\beta)} \tag{5}$$

Eqns. 3-5 constitute all of the terms in Bayes' Rule (eqn. 2). Substituting and solving yields eqn. 6. This equation represents the probability for each possible value for x given k (the number of 1's received) and n (the total number of bits received).

$$p(x|k) = \frac{p(k|x)p(x)}{p(k)}$$
$$= \frac{\frac{n!}{k!(n-k)!}x^k(1-x)^{n-k}\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}x^{\alpha-1}(1-x)^{\beta-1}}{\frac{n!}{k!(n-k)!}\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\frac{\Gamma(k+\alpha)\Gamma(n-k+\beta)}{\Gamma(n+\alpha+\beta)}}$$
$$= \frac{\Gamma(n+\alpha+\beta)}{\Gamma(k+\alpha)\Gamma(n-k+\beta)} \cdot x^{k+\alpha-1}(1-x)^{n-k+\beta-1} \tag{6}$$

The probability density function shown in eqn. 6 is renormalized such that the random variable is not the probability, x, but instead the range of all possible packet values. This is accomplished by a change of variables. v, the value of the transmitted packet, is a function of the random variable X, as shown in eqn. 7.

$$V = MX \tag{7}$$

M in eqn. 7 is the maximum integer that can be stored in the packet length. A change of variables for eqn. 6 yields eqn. 8, the probability density function for all possible transmitted packets, given the number of received 1's (k) and the total number of received bits (n). The probability that the transmitted packet is within a particular range of possible values is calculated by integrating this probability density function over that range of values.

$$p(v|k) = \left[\frac{\Gamma(n+\alpha+\beta)}{\Gamma(k+\alpha)\Gamma(n-k+\beta)} \cdot \left(\frac{v}{M}\right)^{k+\alpha-1} \cdot \right.$$
$$\left. \left(1-\frac{v}{M}\right)^{n-k+\beta-1} \cdot \left(\frac{1}{M}\right)\right] \tag{8}$$

VI. CASE STUDY

Consider the concrete example of a sender equipped with an 8-bit analog-to-digital converter (ADC). The sender must transmit an 8-bit unsigned measurement, v, from the ADC through the deletion channel. It does so by first calculating x, as shown in eqn. 9. In the case of an 8-bit packet, $M = 256$.
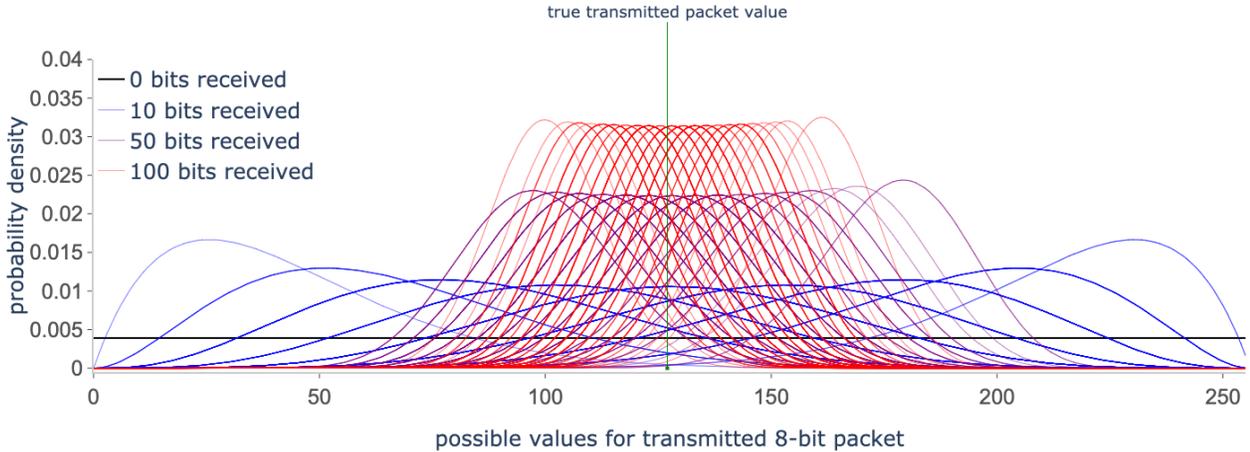
$$x = \frac{v}{M} = \frac{v}{256} \tag{9}$$

Fig. 2. All a posteriori distributions for 100 transmissions of the same packet with 0 bits received, 10 bits received, 50 bits received, and 100 bits received

A random number generator on the sender then generates uniformly distributed random numbers in the range [0,1]. In the case that a randomly generated number is less than x, the sender transmits a 1. In the case that the randomly generated number is greater than x, the sender transmits a 0. Each of these bits may be deleted by the channel with an unknown (and arbitrarily high) deletion probability p. The positions of these deletions are unknown to both the sender and the receiver. After sending N bits (N also unknown to the receiver, and unknown a priori to the sender), the channel closes.

The receiver recovers n ($1 \leq n \leq N$) bits, k of which are 1's. The receiver then uses eqn. 8 to determine the probability density for each possible packet, v, that may have been transmitted given the number of received 1's, k, and the total number of received packets, n. We will assume that all packets are equally likely, and therefore assign hyperparameters α and β both value 1.

Fig. 1 shows a particular instantiation of this random process as a concrete example. In the particular instantiation shown in Fig. 1, the maximum of the a posteriori distribution is near to the true value of the transmitted packet. However, the a posteriori distribution has nonzero probability density for all possible values for the transmitted packet. It is possible, therefore, for the maximum of the a posteriori distribution (which is the same as the maximum likelihood estimate in the case of a uniform prior) to be far from the true value, though this becomes increasingly rare as the number of received bits increases. This is illustrated in Fig. 2, which shows all of the a posteriori distributions for 100 transmissions of the same packet with 0 bits received, 10 bits received, 50 bits received, and 100 bits received. Fig. 2 shows that the maximum of the a posteriori distribution tends toward the true value as the number of bits increases, and that the variance of the distribution decreases. The statistics of this convergence are discussed in Section vII.

VII. THE STATISTICS OF CONVERGENCE

Fig. 2 suggests that, as more bits are received, the expected value of the a posteriori distribution converges to the true transmitted packet value (i.e. the estimator is consistent) and the variance decreases. This section proves those two observations and quantifies the rate of convergence. Section VII.A finds the expected value of the a posteriori distribution as a function of k, α, β, and M. Section VII.B calculates the upper limit for the expected error between the expected value for the a posteriori distribution and the true transmitted packet value, and shows that this error goes to zero as n approaches infinity. Section VII.C calculates the maximum and minimum variance of the a posteriori distribution for a uniform prior distribution.

A. Expected a posteriori estimate

The expected value for the a posteriori distribution is calculated in eqn. 10.

$$
\begin{aligned}
\hat{v}_{EAP} &= \int_v v \cdot p(v|k)dv \\
&= \int_0^M v \cdot \left[\frac{\Gamma(n+\alpha+\beta)}{\Gamma(k+\alpha)\Gamma(n-k+\beta)} \cdot \left(\frac{v}{M}\right)^{k+\alpha-1} \cdot \right. \\
&\qquad \left. \left(1-\frac{v}{M}\right)^{n-k+\beta-1} \cdot \left(\frac{1}{M}\right) \right] dv \\
&= \int_0^1 x \cdot M \cdot \left[\frac{\Gamma(n+\alpha+\beta)}{\Gamma(k+\alpha)\Gamma(n-k+\beta)} \cdot x^{k+\alpha-1} \cdot \right. \\
&\qquad \left. (1-x)^{n-k+\beta-1} \cdot \left(\frac{1}{M}\right) \right] \cdot M dx \\
&= M \cdot \frac{k+\alpha}{n+\alpha+\beta}
\end{aligned} \tag{10}
$$

B. Upper bound on expected error of the expected a posteriori estimate

The error between the expected value for the a posteriori distribution and the true transmitted packet value, $\tilde{v}$ is calculated in eqn. 11.

$$
\begin{aligned}
\tilde{v} &= v - \hat{v}_{EAP} \\
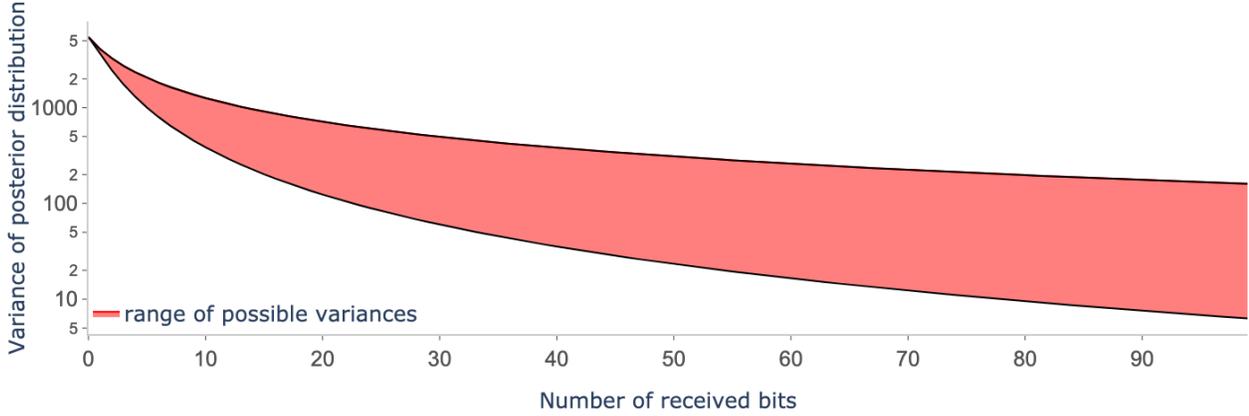&= v - M \cdot \frac{k+\alpha}{n+\alpha+\beta}
\end{aligned} \tag{11}
$$

Fig. 3. Range of possible variances for *a posteriori* distribution as a function of number of received bits, assuming an 8-bit packet.

This error is a random variable, for which the expected value is calculated in eqn. 12. From eqn. 12, it can be seen that the expected value for the error goes to zero as n approaches infinity. Thus, the estimator is consistent.

$$
\begin{aligned}
E[\tilde{v}] &= E\left[v - M \cdot \frac{k+\alpha}{n+\alpha+\beta}\right] \\
&= v - \frac{M}{n+\alpha+\beta} \cdot (E[k]+\alpha) \\
&= v - \frac{M}{n+\alpha+\beta} \cdot (xn+\alpha) \\
&= v - \frac{M}{n+\alpha+\beta} \cdot \left(\frac{v}{M} \cdot n + \alpha\right) \\
&= \frac{v(\alpha+\beta) - M\alpha}{n+\alpha+\beta}
\end{aligned}
\tag{12}
$$

Eqn. 13 is the particular form that eqn. 12 assumes under the assumption of a uniform prior distribution for v.

$$
E[\tilde{v}] = \frac{2v - M}{n+2}
\tag{13}
$$

For a uniform prior distribution ($\alpha = \beta = 1$) and fixed n, the expected error is maximized when v is at the limits of its range (0 or M). In the case of a uniform prior, an upper limit for the expected error can be found by setting $v = M$, as shown in eqn. 14.

$$
|E[\tilde{v}]| \leq \frac{M}{n+2}
\tag{14}
$$

The upper bound for the expected value of the error goes to 0 as n approaches infinity (the estimator is consistent), and is less than 1 (the smallest value separating valid packets) when $n > (M-2)$. This does not guarantee that the expected *a posteriori* estimate for the packet value will be within 1 of the true value when $n > (M-2)$. The posterior distribution has variance about the expected value.

C. Variance of the a posteriori distribution

The variance for the *a posteriori* distribution is calculated as shown in eqn. 15.

$$
\begin{aligned}
\mathrm{Var}[v] &= E\left[v^2\right] - E\left[v\right]^2 \\
&= E\left[v^2\right] - \hat{v}_{EAP}^2
\end{aligned}
\tag{15}
$$

The second term in eqn. 15, $\hat{v}_{EAP}$ is calculated in eqn. 12. The first term, $E[v^2]$, is calculated in eqn. 16.

$$
\begin{aligned}
E\left[v^2\right] &= \int_v v^2 \cdot p(v|k)dv \\
&= \int_0^M v^2 \cdot \left[\frac{\Gamma(n+\alpha+\beta)}{\Gamma(k+\alpha)\Gamma(n-k+\beta)} \cdot \left(\frac{v}{M}\right)^{k+\alpha-1} \cdot \right. \\
&\qquad \left. \left(1-\frac{v}{M}\right)^{n-k+\beta-1} \cdot \left(\frac{1}{M}\right)\right] dv \\
&= M^2 \frac{(k+\alpha+1)(k+\alpha)}{(\alpha+n+\beta+1)(\alpha+n+\beta)}
\end{aligned}
\tag{16}
$$

The variance for v can be calculated by substituting eqn. 12 and eqn. 16 into eqn. 15, as shown in eqn. 17.

$$
\begin{aligned}
\mathrm{Var}[v] &= \left[M^2 \frac{(k+\alpha+1)(k+\alpha)}{(\alpha+n+\beta+1)(\alpha+n+\beta)} - \right. \\
&\qquad \left. \left(M \cdot \frac{k+\alpha}{n+\alpha+\beta}\right)^2\right] \\
&= M^2 \left[\frac{(k+\alpha)(n-k+\beta)]}{(\alpha+n+\beta+1)(\alpha+n+\beta)^2}\right]
\end{aligned}
\tag{17}
$$

In the case of a uniform prior distribution for v (all packets equally likely), then $\alpha = \beta = 1$ and eqn. 17 reduces to eqn. 18.

$$
\mathrm{Var}[v] = M^2 \left[\frac{(k+1)(n-k+1)]}{(n+3)(n+2)^2}\right]
\tag{18}
$$

For a given n, this variance is maximized when $k = \frac{n}{2}$ and minimized when $k = 0$ or $k = n$. Fig 3 shows the range of possible variances of the posterior distribution with increasing number of received bits, n, and in the particular case of an 8-bit packet.
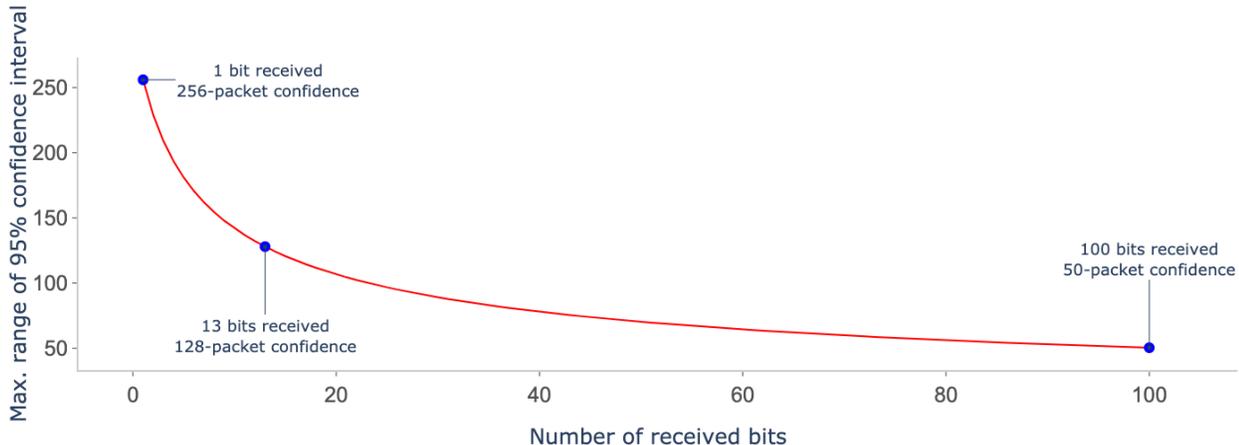
Fig. 4. Upper bound on the number possible packet values contained within the 95 percent confidence interval of the posterior distribution which maximizes variance for the particular case of an 8-bit packet (256 possible packets).

VIII. BOUNDING THE CONFIDENCE INTERVAL OF THE POSTERIOR DISTRIBUTION FOR UNIFORM PRIOR DISTRIBUTION

The posterior distribution, shown in eqn. 8, is a beta distribution that can be rewritten as shown in eqn. 19. From eqn. 8, it can be seen that the effective α and β values for the posterior distribution (denoted α_e and β_e in eqn. 19) are given by the expressions in eqn. 20-21.

$$p(v|k) = \left[\frac{\Gamma(\alpha_e + \beta_e)}{\Gamma(\alpha_e)\Gamma(\beta_e)} \cdot \left(\frac{v}{M}\right)^{\alpha_e - 1} \cdot \right.$$
$$\left. \left(1 - \frac{v}{M}\right)^{\beta_e - 1} \cdot \left(\frac{1}{M}\right) \right] \qquad (19)$$

The variance of this beta distribution is shown in eqn. 18. From eqn. 18, it can be seen that the variance is maximized when $k = \frac{n}{2}$.

$$\alpha_e = k + \alpha \qquad (20)$$
$$\beta_e = n - k + \beta \qquad (21)$$

Under the assumption of a uniform prior distribution, $\alpha = \beta = 1$. Substituting $k = \frac{n}{2}$ yields the effective α_e and β_e shown in eqn. 22-23.

$$\alpha_e = \frac{n}{2} + 1 \qquad (22)$$
$$\beta_e = \frac{n}{2} + 1 \qquad (23)$$

When the α and β values for a beta distribution are equal, they are increasingly well approximated by a Gaussian distribution as their values increase. This is the case for the particular value of k which maximizes the variance, so the maximum confidence intervals are well approximated by treating the posterior distribution with maximum variance as a Gaussian distribution with the same variance.

Fig. 4 shows the number of possible packet values contained within the 95 percent confidence interval of the maximum-variance posterior distribution as a function of the number of received bits. This plot is for the particular case of an

8-bit packet. When only a single bit has been received, the 95 percent confidence interval includes all possible packets. When 13 bits have been received, the 95 percent confidence interval contains 128 of the 256 possible packets. When 100 bits have been received, the 95 percent confidence interval includes approximately 50 of the 256 possible packets. These values empirically agree with Fig. 2.

IX. CONCLUSION

This paper presents and analyzes a method for probabilistically encoding a packet for transmission through a binary deletion channel with an arbitrarily high deletion probability and unknown channel lifetime. The number of bits received before channel closure is unknown to both sender and receiver and assumed greater than or equal to one bit and less than or equal to 100 bits. The supposed deletion channel is one with properties that do not guarantee the intact receipt of any packet larger than one bit, and that do not guarantee the receipt of enough bits to perfectly reconstruct packets longer than 1 bit from partial packets. For such channels, this paper argues that the appropriate encoding method is one which enables the calculation of probability masses over all possible packets rather than one which attempts perfect reconstruction of the transmitted packet. This paper presents such an encoding method.

This paper proposes to encode information in the probability of receiving a 1 vs. a 0 (rather than in the 1's and 0's themselves). Assuming a fixed packet length, the packet to be transmitted is interpreted as an integer. The ratio between this integer and the largest integer storable in the packet length is interpreted as a probability, x. The sender uses x to decide whether to transmit a binary 1 or 0. Each received bit is modeled as an independent and identically distributed Bernoulli random trial. The bit is 1 with probability x and 0 with probability $(1-x)$. The probability density function for x is calculated by the receiver via Bayesian inference. This paper analyzes the statistics of convergence of this estimator and finds the upper bound for confidence intervals for all numbers

of received bits from 0 to 100. It is shown that, for an 8-bit packet with 256 valid integer values, the upper bound for the 95 percent confidence interval contains 128 of the 256 possible packets when 13 bits have been received. When 100 bits have been received, the 95 percent confidence interval includes approximately 50 of the 256 possible packets.

ACKNOWLEDGMENT

Thank you to Kalani Danas Rivera, Joshua Umansky-Castro, and Rachel Miller for reviewing this work. [7] was used as a reference for the standard statistics equations used throughout this paper.

REFERENCES

[1] "Starshot." Breakthrough Initiatives, Breakthrough Initiatives, break-throughinitiatives.org/initiative/3.

[2] Levenshtein, Vladimir I. "Binary codes capable of correcting deletions, insertions, and reversals." Soviet physics doklady. Vol. 10. No. 8. 1966.

[3] Diggavi, Suhas, Michael Mitzenmacher, and Henry D. Pfister. "Capacity upper bounds for the deletion channel." 2007 IEEE International Symposium on Information Theory. IEEE, 2007.

[4] Fertonani, Dario, and Tolga M. Duman. "Novel bounds on the capacity of the binary deletion channel." IEEE Transactions on Information Theory 56.6 (2010): 2753-2765.

[5] Drinea, Eleni, and Michael Mitzenmacher. "Improved lower bounds for the capacity of iid deletion and duplication channels." IEEE Transactions on Information Theory 53.8 (2007): 2693-2714.

[6] Y. Kanoria and A. Montanari, "Optimal Coding for the Binary Deletion Channel With Small Deletion Probability," in IEEE Transactions on Information Theory, vol. 59, no. 10, pp. 6192-6219, Oct. 2013.

[7] Bar-Shalom, Y., Li, X. R., & Kirubarajan, T. (2004). Estimation with applications to tracking and navigation: theory algorithms and software. John Wiley & Sons.