Timely Updates over an Erasure Channel

†Roy D. Yates, ‡Elie Najm, †Emina Soljanin, and †Jing Zhong,
† ECE Dept., Rutgers University, {ryates, emina.soljanin, jing.zhong}@rutgers.edu, ‡LTHI, EPFL, elie.najm@epfl.ch

***Abstract*—Using an age of information (AoI) metric, we examine the transmission of coded updates through a binary erasure channel to a monitor/receiver. We start by deriving the average status update age of an infinite incremental redundancy (IIR) system in which the transmission of a k-symbol update continues until k symbols are received. This system is then compared to a fixed redundancy (FR) system in which each update is transmitted as an n symbol packet and the packet is successfully received if and only if at least k symbols are received. If fewer than k symbols are received, the update is discarded. Unlike the IIR system, the FR system requires no feedback from the receiver. For a single monitor system, we show that tuning the redundancy to the symbol erasure rate enables the FR system to perform as well as the IIR system. As the number of monitors is increased, the FR system outperforms the IIR system that guarantees delivery of all updates to all monitors.**

I. Introduction

Status update systems have focused on sending updates to a monitor through a system or network in which the transmission of an update requires a random service time [1]–[6]. In this work, we consider a system in which a source sends coded updates through an unreliable channel to a monitor. We examine how to employ coding redundancy in order to minimize an "Age of Information" (AoI) metric. We will see that this is not the same as ensuring reliable delivery of every update while minimizing the coding overhead. Over a noisy channel, the age is reduced only when an update is correctly received, but using coding to increase the probability of correct reception incurs additional delay for each update. In this work, we examine what is just the right amount of redundancy.

In our system model, the source sends updates through a symbol erasure channel to a monitor. One symbol is transmitted per unit time. A symbol is erased with probability δ; otherwise it is received correctly. Each update is a timestamped file with k information symbols. Depending on the coding strategy, these updates are encoded into at least k and possibly infinitely many (rateless) coded symbols for transmission over the channel. The monitor may employ a feedback channel to notify the source about symbols that have been erased or packets that have failed to be decoded.

If at time t, the most recent received update is timestamped $u(t)$, the status age is $\Delta(t) = t - u(t)$. In the absence of an update, the status age increases linearly with time. Thus the age process $\Delta(t)$ is a sawtooth waveform as shown in Figure 1. To compare coding strategies, our system performance metric is the time-average status age (also known as the AoI)

$$\Delta = \lim_{T\to\infty} \frac{1}{T}\int_0^T \Delta(t)\,dt. \tag{1}$$

We will use two coding techniques: 1) an infinite incremental redundancy (IIR) strategy and 2) a finite redundancy (FR) strategy. Under the IIR strategy, each k-symbol update is encoded by a rateless code such that when k coded symbols are correctly received by the monitor, the update is successfully decoded (e.g., a Reed-Solomon or a Fountain code). The source is provided instantaneous feedback when the update has been decoded, at which point it starts transmitting a new update.

Under the FR strategy, each k-symbol update is encoded as an n-symbol packet. The update is successfully delivered as soon as k un-erased symbols are received. If fewer than k symbols are received, the update is discarded. This system employs no feedback from the monitor and thus all n symbols of an update are transmitted even if the monitor successfully decodes the update before the transmission is finished. The source starts transmitting a new update once the n symbols of the previous update have been sent. Note that under the FR strategy 1) not every update will be decoded and 2) there generally will be a positive time gap between the completion of decoding of an update and the beginning of the new update transmission.

In this work, Section II analyzes AoI for the IIR system, first with a single monitor and then with $m > 1$ monitors. In Section III, we characterize age in the FR updating system. For this system, we show that by matching the redundancy n to the erasure rate, the FR system has AoI approaching that of the single-monitor IIR system as k becomes large. In Section IV, we present numerical evaluations of both systems. A brief discussion concludes this work in Section V. The page limit precluded the inclusion of the Appendix. A version with this appendix is available [7].

II. AoI Under the IIR Strategy

A. Single Monitor System

Update 1 begins transmission at time $t = 0$ and is timestamped $T_0 = 0$. To analyze the average age, we define X_i as the number of symbols sent until the kth un-erased symbol of update i is received. Because the erasure channel is memoryless, $X_1, X_2, \ldots$ are iid negative binomial (NB) $(k, 1-\delta)$ random variables, identical to $X^{(k)}$ with PMF

$$P_{X^{(k)}}(x) = \binom{x-1}{k-1}(1-\delta)^k\delta^{x-k}, \quad x = k, k+1, \ldots \tag{2}$$

For convenience, we will denote the CDF of $X^{(k)}$ by

$$F_k(n) = \sum_{x=k}^{n}\binom{x-1}{k-1}(1-\delta)^k\delta^{x-k}, \ n = k, k+1, \ldots \tag{3}$$

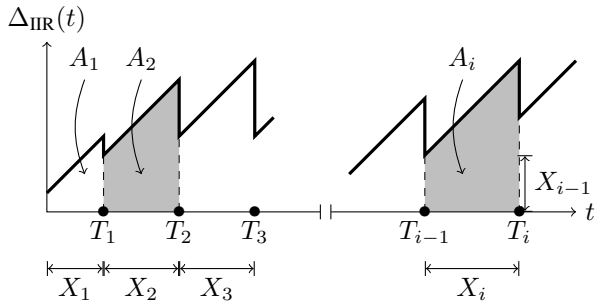

Fig. 1. Sample path of the status update age $\Delta_{\text{IIR}}(t)$ (the upper envelope in bold) for the IIR updating system. Updates are delivered at time instances marked by •. Update i submitted at time T_{i-1} is delivered after a transmission time X_i.

We also note that $X^{(k)}$ has expected value $\mathrm{E}[X^{(k)}] = \mu_k$ and variance $\mathrm{Var}[X^{(k)}] = \sigma_k^2$ with

$$\mu_k = \frac{k}{1-\delta}, \qquad \sigma_k^2 = \frac{k\delta}{(1-\delta)^2}. \tag{4}$$

Following the delivery of update l at time $T_l = \sum_{i=1}^{l} X_l$, update $l+1$ immediately begins transmission.

To analyze the average age Δ, we decompose the area defined by the integral (1) into a sum of disjoint polygonal areas $A_1, A_2, \ldots$ as shown in Figure 1. Over the time interval $(0, T = T_l)$, this decomposition yields average age

$$\Delta_{\text{IIR}} = \lim_{l\to\infty} \frac{1}{T_l} \sum_{i=1}^{l} A_i = \lim_{l\to\infty} \frac{\frac{1}{l}\sum_{i=1}^{l} A_i}{\frac{1}{l}\sum_{i=1}^{l} X_i} = \frac{\mathrm{E}[A]}{\mathrm{E}[X]}. \tag{5}$$

When update i begins transmission at time T_{i-1}, the age is $\Delta(T_{i-1}) = X_{i-1}$. From Figure 1, we see that the area A_i is

$$A_i = X_{i-1}X_i + X_i^2/2. \tag{6}$$

Since the X_i are iid, $\mathrm{E}[A] = (\mathrm{E}[X])^2 + \mathrm{E}\big[X^2\big]/2$ and it follows from (4) and (5) that the average age of the IIR system is

$$\Delta_{\text{IIR}} = \mathrm{E}[X] + \frac{\mathrm{E}\big[X^2\big]}{2\,\mathrm{E}[X]} = \frac{k}{1-\delta}\left(\frac{3}{2} + \frac{\delta}{k}\right). \tag{7}$$

We note IIR is the only strategy that guarantees the delivery of every update. Moreover, it minimizes the coding overhead, and thus maximizes the throughput. It takes $k/(1-\delta)$ coded symbols on average to transmit a k-symbol update, which is not equal to the average update age. In particular, $3k/[2(1-\delta)]$ is what the average age would be if each update were delivered by exactly $k/(1-\delta)$ symbol transmissions. The additional (though admittedly small) age penalty of IIR reflects the randomness in the negative binomial distribution.

We also observe that IIR is a zero-wait system: as soon as an update is delivered, a new update goes into service. However, when service times are random, zero-wait policies may not be age-minimizing. By [6, Theorem 5], it can be shown that zero-wait is optimal for IIR if and only if $\delta \le k/(2k+1)$.

B. Multiple Monitor System

Using IIR to transmit to $m > 1$ monitors, the source continues to transmit encoded symbols until each of the m monitors has correctly received k coded symbols. The source is provided instantaneous feedback when an update has been decoded by all users.

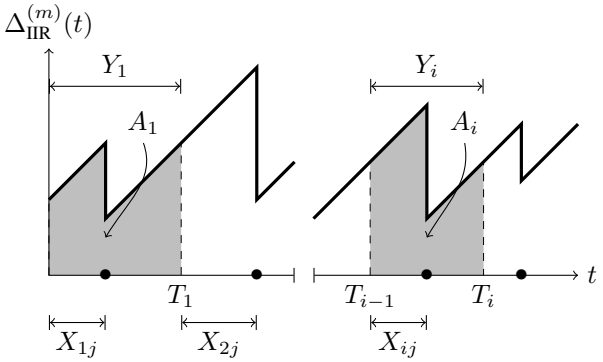

Fig. 2. Sample path of the status update age $\Delta_{\text{IIR}}^{(m)}(t)$ for user j in the IIR updating system with $m > 1$ monitors. Update i completes transmission at time T_i. Update delivery instances for monitor j are marked by •.

Update 1 begins transmission at time $t = 0$ and is timestamped $T_0 = 0$. To analyze the average age, we define X_{ij} as the number of symbols sent until the kth un-erased symbol of update i is received by monitor j. Because the erasure channels of all users are memoryless and independent, the X_{ij} are iid NB $(k, 1-\delta)$ random variables with PMF given by (2).

The transmission time of update i is

$$Y_i = \max(X_{i1}, \ldots, X_{im}). \tag{8}$$

The Y_i are an iid sequence, each with CDF

$$F_Y(y) = \mathrm{P}[Y \le y] = \mathrm{P}[X^{(k)} \le y]^m = [F_k(y)]^m. \tag{9}$$

Following the delivery of update l at time $T_l = \sum_{i=1}^{l} Y_i$, update $l+1$ immediately begins transmission.

Since all monitors have statistically identical (but independent) channels, we define $\Delta(t)$ as the age of some monitor j and we now analyze the average age Δ. Figure 2 depicts the age process $\Delta(t)$ for monitor j. The analysis of the average age is similar to that for the single user IIR system. As before, update i completes transmission at time T_i, but, for user j, the age $\Delta(t)$ drops when update i is delivered to monitor j at the earlier time $T_{i-1} + X_{ij}$. We note this implies that update i completes transmission at time T_i, the age at monitor j is then $\Delta(T_i) = Y_i$.

As we did for the single user system, we represent the area of the integral (1) as the concatenation of the polygons $A_1, \ldots, A_l$, yielding the average age $\Delta_{\text{IIR}}^{(m)} = \mathrm{E}[A]/\,\mathrm{E}[Y]$. Examination of Figure 2 will show that

$$\begin{aligned} A_i &= Y_{i-1}X_{ij} + X_{ij}^2/2 + X_{ij}(Y_i - X_{ij}) + (Y_i - X_{ij})^2/2 \\ &= Y_{i-1}X_{ij} + Y_i^2/2. \end{aligned} \tag{10}$$

Since X_{ij} is independent of the transmission time Y_{i-1} of the

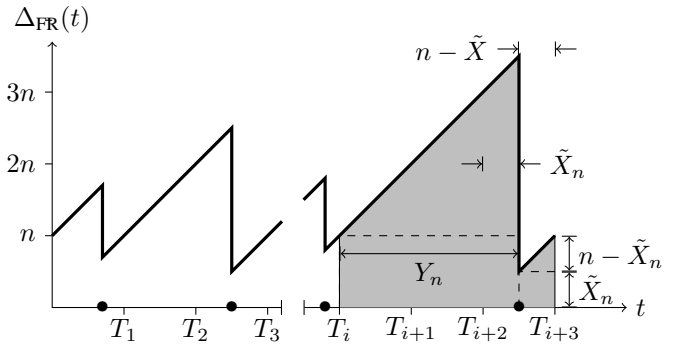

Fig. 3. A sample path of the FR age $\Delta_{\text{FR}}(t)$: successful update deliveries (at times marked by $\bullet$) occur in slots 1, 3, i, and $i+3$. Updates are discarded in slots 2, $i+1$, and $i+2$.

previous update, $\mathrm{E}[A] = \mathrm{E}[Y]\,\mathrm{E}[X] + \mathrm{E}\big[Y^2\big]/2$ and

$$\Delta_{\text{IIR}}^{(m)} = \mathrm{E}[X] + \frac{\mathrm{E}\big[Y^2\big]}{2\,\mathrm{E}[Y]}. \tag{11}$$

Using (3) and (9), the moments $\mathrm{E}[Y]$ and $\mathrm{E}\big[Y^2\big]$ are easy to calculate but they do not have simple closed form expressions.

III. AoI under Fixed Redundancy Coding

Under the fixed redundancy (FR) strategy, each update is encoded as an n-symbol packet but the update is successfully decoded as soon as k un-erased symbols are received. If fewer than k symbols are received, the update is discarded. This system employs no feedback from the monitor and thus all n symbols of an update are transmitted even if the monitor decodes the update before the transmission is finished.

To analyze this system, we define $T_i = in$ and slot i as the time interval $(T_{i-1}, T_i]$. Update i is successfully delivered in slot i with probability

$$1 - \epsilon_n = \mathrm{P}\Big[X^{(k)} \le n\Big] = F_k(n). \tag{12}$$

Referring to Figure 3, when a success occurs in slot i, the age at time $T_{i-1} + X_i$ is reset to X_i because update i has age X_i at that time instant. Moreover, $\Delta(T_i) = n$ because update i will then have age n at the end of slot i. Consequently, when a success occurs in slot i, T_i is a renewal point of the process $\Delta(t)$ in that $\Delta(T_i) = n$ and time instant T_i marks the start of transmission of a fresh update. In the example of Figure 3, renewals occur at times T_1, T_3, T_i, and T_{i+3}.

Measured in slots, the length of a renewal period is a geometric $(1-\epsilon_n)$ random variable M_n with PMF

$$P_{M_n}(m) = \epsilon_n^{m-1}(1-\epsilon_n), \qquad m = 1, 2, \ldots, \tag{13}$$

corresponding to $m-1$ updates being discarded followed by a success with update m.

For the FR system, we analyze the AoI Δ in (1) using renewal-reward theory [8]. Specifically, we interpret $\Delta(t)$ as an instantaneous reward rate so that Δ is the average reward rate. In the renewal period starting at time T_i, the reward

$$R = \int_{T_i}^{T_{i+M_n}} \Delta(t)\,dt \tag{14}$$

is earned. In Figure 3, R is the shaded area. This renewal period terminates after $M_n = m$ slots because $X_{i+m} \le n$. This implies X_{i+M_n} is identical to a random variable $\tilde{X}_n$ with PMF $P_{\tilde{X}_n}(x) = P_{X^{(k)}|X^{(k)}\le n}(x)$. From (2) and (3),

$$P_{\tilde{X}_n}(x) = \frac{\binom{x-1}{k-1}(1-\delta)^k\delta^{x-k}}{F_k(n)}, \quad x \le n. \tag{15}$$

It will be convenient to define $\tilde{\mu}_n \equiv \mathrm{E}[\tilde{X}_n]$ and we note that (15) implies

$$\begin{aligned}\tilde{\mu}_n &= \frac{1}{F_k(n)}\sum_{x=k}^{n} x\binom{x-1}{k-1}(1-\delta)^k\delta^{x-k} \\ &= \frac{k}{F_k(n)}\sum_{x=k}^{n}\binom{x}{k}(1-\delta)^k\delta^{x-k}.\end{aligned} \tag{16}$$

With the substitutions $x' = x+1$ and $k' = k+1$, we obtain

$$\begin{aligned}\tilde{\mu}_n &= \frac{k}{(1-\delta)F_k(n)}\sum_{x'=k'}^{n+1}\binom{x'-1}{k'-1}(1-\delta)^{k'}\delta^{x'-k'} \\ &= \frac{kF_{k+1}(n+1)}{(1-\delta)F_k(n)}.\end{aligned} \tag{17}$$

Note that $\tilde{X}_n$ is independent of the number of slots M_n in a given renewal period. Referring to Figure 3, the renewal period consists of an interval of length

$$Y_n = n(M_n - 1) + \tilde{X}_n = nM_n - (n - \tilde{X}_n). \tag{18}$$

in which $\Delta(t)$ grows from $\Delta(T_j) = n$ to $\Delta(T_j + Y) = n + Y_n$, followed by a second interval of length $n - \tilde{X}_n$. As shown in the figure, each of these intervals contributes a rectangular area and a triangular area to the reward R. Thus,

$$\begin{aligned}R &= nY_n + Y_n^2/2 + \tilde{X}_n(n - \tilde{X}_n) + (n - \tilde{X}_n)^2/2 \\ &= nY_n + Y_n^2/2 + n^2/2 - \tilde{X}_n^2/2.\end{aligned} \tag{19}$$

It then follows from (18) that

$$R = n^2M_n^2/2 + nM_n\tilde{X}_n. \tag{20}$$

Since the renewal period has length $M_n n$, the renewal-reward theorem ensures that the time-average reward (corresponding to the time-average age Δ) is

$$\Delta_{\text{FR}}(n) = \frac{\mathrm{E}[R]}{\mathrm{E}[M_n n]} = \frac{n\,\mathrm{E}\big[M_n^2\big]}{2\,\mathrm{E}[M_n]} + \tilde{\mu}_n. \tag{21}$$

Since M_n has moments

$$\mathrm{E}[M_n] = \frac{1}{1-\epsilon_n}, \qquad \mathrm{E}\big[M_n^2\big] = \frac{1+\epsilon_n}{(1-\epsilon_n)^2}, \tag{22}$$

it follows that

$$\Delta_{\text{FR}}(n) = \frac{n}{1-\epsilon_n} - \frac{n}{2} + \tilde{\mu}_n. \tag{23}$$

We note that calculation of $\Delta_{\text{FR}}(n)$ is straightforward using

(12) and (17).

A. AoI Bounds under FR

We will see from numerical evaluations in Section IV that given k, δ there exists an optimal redundancy n_k^* such that

$$\Delta_{\text{FR}}^* = \Delta_{\text{FR}}(n_k^*) \le \Delta_{\text{FR}}(n) \tag{24}$$

for all n. To characterize n_k^*, we now derive $\Delta_{\overline{\text{FR}}}(n)$, a surprisingly tight upper bound on the average age $\Delta_{\text{FR}}(n)$. We then show that a close approximation to n_k^* can be found by a minimization of $\Delta_{\overline{\text{FR}}}(n)$ based on a central limit theorem (CLT) approximation. While this method is approximate, the result will yield a strict (and tight) upper bound to Δ_{FR}^*. We start with the following claim, with proof in the Appendix.

Lemma 1: For fixed k and δ, the sequence $\tilde{\mu}_k, \tilde{\mu}_{k+1}, \ldots$ is nondecreasing and satisfies

$$\tilde{\mu}_n \le \min(n, k/(1-\delta)).$$

Applying Lemma 1 to (23), we obtain the upper bound

$$\Delta_{\text{FR}}(n) \le \Delta_{\overline{\text{FR}}}(n) \equiv \frac{n}{1-\epsilon_n} - \frac{n}{2} + \frac{k}{1-\delta}. \tag{25}$$

Writing $n = \sigma_k z + \mu_k$, we employ the CLT approximation

$$\begin{aligned} 1-\epsilon_n &= \text{P}\left[X^{(k)} \le n\right] \\ &= \text{P}\left[X^{(k)} \le \sigma_k z + \mu_k\right] \approx \Phi(z) \end{aligned} \tag{26}$$

where $\Phi(z)$ is the standard Gaussian CDF. Applied to (25), this approximation permits us to write

$$\begin{aligned} \Delta_{\overline{\text{FR}}}(n) &\approx \frac{\sigma_k z + \mu_k}{\Phi(z)} - \frac{\sigma_k z + \mu_k}{2} + \mu_k \\ &= \sigma_k \left(\frac{z + \hat{\mu}_k}{\Phi(z)} - \frac{z}{2} + \frac{\hat{\mu}_k}{2} \right) \end{aligned} \tag{27}$$

where $\hat{\mu}_k \equiv \mu_k/\sigma_k = \sqrt{k/\delta}$. Because $\hat{\mu}_k \gg z$ for typical parameters, we make the further approximation

$$\Delta_{\overline{\text{FR}}}(n) \approx \sigma_k \left(\frac{\hat{\mu}_k}{\Phi(z)} - \frac{z}{2} + \frac{\hat{\mu}_k}{2} \right). \tag{28}$$

Setting the derivative of the right side of (28) to zero, we obtain $-\hat{\mu}_k \Phi'(z) = [\Phi(z)]^2/2$. For large k, we will want the probability an update is decoded to be fairly close to 1. Hence $\Phi(z) \approx 1$ for values of z of interest. Since $\Phi'(z) = e^{-z^2/2}/\sqrt{2\pi}$, solving $\Phi'(z) = -1/[2\hat{\mu}_k]$ yields $z = z_k^* = \sqrt{\ln(2k/\pi\delta)}$. Employing (4), we obtain the threshold

$$\hat{n}_k^* = \mu_k + \sigma_k z_k^* = \frac{k}{1-\delta}(1+\omega_k) \tag{29}$$

where

$$\omega_k \equiv \left[\frac{\delta}{k} \ln \frac{2k}{\pi\delta} \right]^{1/2}. \tag{30}$$

In the Appendix, we verify the following claim:

Lemma 2: Given $\eta_0 > 0$, there exists K_0 such that

$$\text{P}\left[X^{(k)} > \hat{n}_k^*\right] \le \beta_k \equiv e^{\eta_0/(1-\delta)} \sqrt{\frac{\pi\delta}{2k}}, \qquad k \ge K_0.$$

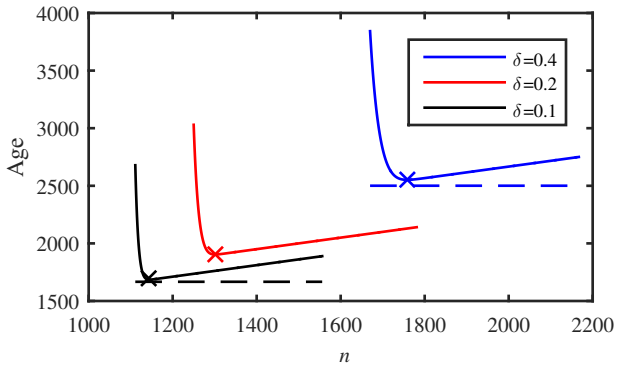

Fig. 4. The average age $\Delta_{\text{FR}}(n)$ for update packets with $k = 1000$ information symbols as a function of the number of transmitted symbols n. For each curve, $\times$ marks $\Delta_{\overline{\text{FR}}}(\hat{n}_k^*)$. For $\delta = 0.4$ and $\delta = 0.1$, the dashed lines show the IIR age Δ_{IIR}.

We note that the tail probability $[\text{P}\left[X^{(k)} > \hat{n}_k^*\right]$ decays slowly (i.e. sub-exponentially) because $\hat{n}_k^*$ is approaching μ_k as k becomes large. It follows from (24), (25) and Lemma 2 that

$$\begin{aligned} \Delta_{\text{FR}}^* \le \Delta_{\overline{\text{FR}}}(\hat{n}_k^*) &= \frac{\hat{n}_k^*}{\text{P}\left[X^{(k)} \le \hat{n}_k^*\right]} - \frac{\hat{n}_k^*}{2} + \frac{k}{1-\delta} \\ &= \frac{k}{1-\delta}\left[\frac{3}{2} + \frac{\beta_k + \frac{1}{2}\omega_k(1+\beta_k)}{1-\beta_k}\right]. \end{aligned} \tag{31}$$

Since β_k and ω_k approach zero as k grows, we see from (7) and (31) that the average age of the IIR system and the average age of FR system with optimized redundancy both asymptotically approach $1.5k/(1-\delta)$.

IV. Evaluation

Figure 4 evaluates a system in which updates have $k = 1000$ information symbols. We plot the FR age $\Delta_{\text{FR}}(n)$ in (23) as a function of n, the FR update packet length, for a range of values of the erasure probability δ. As one would expect, the age increases as δ increases. We also see that for a given erasure probability δ, the optimal n is sharply defined. Too few transmitted symbols and the age blows up because the packet update erasure probability is high; on the other hand, more than the minimum number of sent symbols also creates unnecessary age. Marked by $\times$ are the approximately optimal redundancy $\hat{n}_k^*$ and the corresponding age upper bound $\Delta_{\overline{\text{FR}}}(\hat{n}_k^*)$.

Figure 5 is similar to Figure 4 except there are only $k = 50$ information bits and the figure includes the upper bound $\Delta_{\overline{\text{FR}}}(n)$ in (25). Furthermore, the figure is plotted as the normalized age $\Delta/[k/(1-\delta)]$ vs. the normalized packet length $n/[k/(1-\delta)]$. This normalization and the small value of k are chosen to accentuate the gap between $\Delta_{\text{FR}}(n)$ and $\Delta_{\overline{\text{FR}}}(n)$. For typical values of k such as $k = 1000$, the gap between the age and the upper bound cannot be visually resolved.

Figure 6 compares systems with $m > 1$ monitors. For each value of δ, we compare the IIR age $\Delta_{\text{IIR}}^{(m)}$ and the FR age $\Delta_{\text{FR}}(\hat{n}_k^*)$ using $\hat{n}_k^*$ symbols matched to the erasure rate δ. The IIR system completes the transmission of an update only after all m monitors have decoded. Consequently, the IIR average age grows monotonically with m. Because the FR system ignores whether a monitor has actually decoded an update,

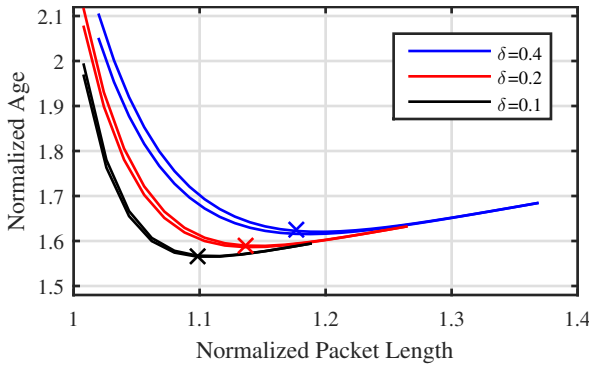

Fig. 5. The average age $\Delta_{\text{FR}}(n)$ and the upper bound $\Delta_{\overline{\text{FR}}}(n)$ for update packets with $k = 50$ information symbols as a function of the number of transmitted symbols n. The age and packet length are normalized by $k/(1-\delta)$. For each curve, $\times$ marks $\Delta_{\overline{\text{FR}}}(\hat{n}_k^*)$.

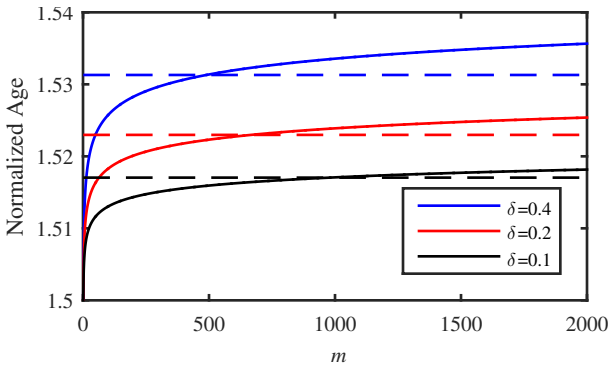

Fig. 6. The average age $\Delta_{\text{IIR}}^{(m)}$ (solid line) and the upper bound $\Delta_{\overline{\text{FR}}}(\hat{n}_k^*)$ (dashed line) for update packets with $k = 1000$ information symbols as a function of the number of monitors m. Note that ages are normalized by $k/(1-\delta)$.

the FR age is insensitive to the number of users. To highlight small differences, all ages are normalized by $k/(1-\delta)$. We see that for all values of the erasure probability δ, the FR system outperforms the IIR when the number of monitors m becomes sufficiently large. We also see that normalized system performance is very similar across a range of erasure rates.

V. Discussion

We have shown that the FR system, which requires no feedback, can essentially match the performance of the IIR system that does require update delivery feedback from the monitor. However, the FR system does require the redundancy to be carefully optimized in response to the channel erasure rate. In practical systems, the erasure rate will vary with time and cannot be assumed to be known. Hence, the FR system will also require some form of feedback to establish the appropriate redundancy level. In practice, systems issues, such as whether receiver feedback can be supported, will determine which approach is better in a particular setting.

In addition, other coded redundancy mechanisms merit examination. For example, the finite incremental redundancy (FIR) strategy [9], just like FR, uses a fixed rate code, but the source is provided instantaneous feedback if the update has been decoded before all its n symbols have been transmitted, at which point it starts transmitting a new update, as in IIR. On the other hand, FIR shares the advantage of FR that updates that were unlucky in transmission can be terminated without waiting for successful decoding.

It is natural to compare timely update delivery with HARQ-aided content download. Hybrid ARQ (HARQ) is a special transmission scheme that combines the conventional ARQ with error correction (see e.g. [10]). Incremental redundancy HARQ (IR-HARQ) schemes adapt their error correcting code redundancy to varying channel conditions, and thus achieve better throughput efficiency than ordinary ARQ. In content download, all content needs to be delivered, and thus these systems have to have a rateless transmission at some level (e.g., conventional ARQ or Fountain codes at the packet level) which will continue until each packet is successfully delivered. For example, in eMBMS, an FR strategy on the physical layer would have a Fountain code at the packet level [11]. Furthermore, content download systems strive to minimize the download time, which, as we have seen in Sec II, is not equivalent to minimizing AoI. The behavior of both systems in a multi-user scenario is similar because of an underlying order statistics phenomenon. Roughly speaking, when there are many users, it is very likely that it will take a long time for some to decode, and putting limits on that time as the FR strategy does, will have an advantage. Content download systems will then have to supplement such systems with an outer rateless code. It would be interesting to compare update delivery with content streaming where all packets have to be delivered in a timely manner.

References

[1] S. Kaul, R. D. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. INFOCOM*, 2012.

[2] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Trans. Info Theory*, vol. 62, no. 4, pp. 1897–1910, April 2016.

[3] C. Kam, S. Kompella, G. D. Nguyen, and A. Ephremides, "Effect of message transmission path diversity on status age," *IEEE Trans. Info Theory*, vol. 62, no. 3, pp. 1360–1374, March 2016.

[4] R. D. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *Proc. IEEE Int'l. Symp. Info. Theory*, 2015.

[5] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," in *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, April 2016, pp. 1–9.

[6] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *CoRR*, vol. abs/1601.02284, 2016, submitted to IEEE Trans. Info. Theory. [Online]. Available: http://arxiv.org/abs/1601.02284

[7] R. D. Yates, E. Najm, E. Soljanin, and Zhong Jing, "Timely Updates over an Erasure Channel," Tech. Rep., 2017. [Online]. Available: http://infoscience.epfl.ch/record/224673

[8] R. G. Gallager, *Stochastic processes: theory for applications*. Cambridge University Press, 2013.

[9] M. Heindlmaier and E. Soljanin, "Isn't hybrid ARQ sufficient?" in *Communication, Control, and Computing (Allerton), 52nd Annual Allerton Conference on*. IEEE, 2014, pp. 563–568.

[10] E. Soljanin, R. Liu, and P. Spasojevic, "Hybrid ARQ with random thansmission assignments," in *Advances in Network Information Theory, Proceedings of a DIMACS Workshop, Piscataway, New Jersey, USA, March 17-19, 2003*, 2003, pp. 321–334.

[11] A. Shokrollahi, M. Luby *et al.*, "Raptor codes," *Foundations and Trends® in Communications and Information Theory*, vol. 6, no. 3–4, pp. 213–322, 2011.

APPENDIX

Proof: (Lemma 1) With the shorthand definitions

$$q_n \equiv \frac{F_k(n)}{(1-\delta)^k}, \qquad \hat{q}_n \equiv \frac{F_{k+1}(n+1)}{(1-\delta)^{k+1}}, \tag{32}$$

we observe that (17) permits us to write

$$\tilde{\mu}_n = k\hat{q}_n/q_n. \tag{33}$$

It follows from (3) and (32) that $q_k = \hat{q}_k = 1$ and that

$$q_{k+1} = 1 + k\delta, \qquad \hat{q}_{k+1} = 1 + (k+1)\delta. \tag{34}$$

These facts imply $\tilde{\mu}_k = k$ and

$$\tilde{\mu}_{k+1} = \frac{k\hat{q}_{k+1}}{q_{k+1}} = k + \frac{k\delta}{1+k\delta}. \tag{35}$$

Thus $\tilde{\mu}_k \le \tilde{\mu}_{k+1}$ and $\tilde{\mu}_{k+1} \le k+1$. We now prove by induction that the sequence $\tilde{\mu}_n$ is nondecreasing and satisfies $\tilde{\mu}_n \le n$. Suppose $\tilde{\mu}_k \le \tilde{\mu}_{k+1} \le \cdots \le \tilde{\mu}_{n-1}$ and that $\tilde{\mu}_i \le i$ for $i < n$. Defining $\gamma_n \equiv \binom{n-1}{k-1}\delta^{n-k}$, it follows from (3) that

$$q_n = q_{n-1} + \gamma_n, \tag{36}$$

$$\hat{q}_n = \hat{q}_{n-1} + \binom{n}{k}\delta^{n-k} = \hat{q}_{n-1} + \frac{n}{k}\gamma_n. \tag{37}$$

This implies

$$\tilde{\mu}_n = \frac{k\hat{q}_n}{q_n} = \frac{k\hat{q}_{n-1} + n\gamma_n}{q_{n-1} + \gamma_n}. \tag{38}$$

By our induction hypothesis, $\tilde{\mu}_{n-1} \le n$, or, equivalently, $k\hat{q}_{n-1} \le nq_{n-1}$. Applying this upper bound to the numerator in (38) yields $\tilde{\mu}_n \le n$. We now observe that (38) and $n \ge \tilde{\mu}_{n-1} = k\hat{q}_{n-1}/q_{n-1}$ also imply

$$\tilde{\mu}_n \ge \frac{k\hat{q}_{n-1} + (k\hat{q}_{n-1}/q_{n-1})\gamma_n}{q_{n-1} + \gamma_n} = \frac{k\hat{q}_{n-1}}{q_{n-1}} = \tilde{\mu}_{n-1}. \tag{39}$$

Finally, we observe from (12) that

$$\lim_{n\to\infty} q_n = \frac{1}{(1-\delta)^k}, \qquad \lim_{n\to\infty} \hat{q}_n = \frac{1}{(1-\delta)^{k+1}}. \tag{40}$$

This implies $\lim_{n\to\infty} \tilde{\mu}_n = k/(1-\delta)$. Since $\tilde{\mu}_n$ is nondecreasing, $\tilde{\mu}_n \le k/(1-\delta)$ for all $n \ge k$. ■

Proof: (Lemma 2) Random variable $X^{(k)}$ has moment generating function $\phi_{X^{(k)}}(s) = [(1-\delta)e^s/(1-\delta e^s)]^k$. By the Chernoff bound, $\ln \mathrm{P}\left[X^{(k)} \ge \hat{n}_k^*\right] \le \min_{s\ge 0} P^{(k)}(s)$ where

$$P^{(k)}(s) = \ln[e^{-s\hat{n}^*}\phi_{X^{(k)}}(s)] \tag{41}$$

$$= k\left[\ln(1-\delta) - s\left(\frac{\delta+\omega_k}{1-\delta}\right) - \ln(1-\delta e^s)\right]. \tag{42}$$

It is straightforward to show that $P_k(s)$ is minimized at

$$s^* = \ln[(1+\omega_k/\delta)/(1+\omega_k)]. \tag{43}$$

Using the shorthand notation $L(x) = \ln(1+x)$, it follows from (43) that

$$P^{(k)}(s^*) = \frac{-k[(\delta+\omega_k)L(\omega_k/\delta) - (1+\omega_k)L(\omega_k)]}{1-\delta}. \tag{44}$$

Defining

$$y_1(k) \equiv k[\delta L(\omega_k/\delta) - L(\omega_k)], \tag{45}$$

$$y_2(k) \equiv k\omega_k[L(\omega_k/\delta) - L(\omega_k)], \tag{46}$$

we observe that

$$P^{(k)}(s^*) = -\frac{y_1(k) + y_2(k)}{1-\delta}. \tag{47}$$

With the definition

$$\ell_k \equiv \ln\left(\frac{2k}{\pi\delta}\right) = \ln k + \ln\left(\frac{2}{\pi\delta}\right), \tag{48}$$

we observe from (30) that $\omega_k^2 = \delta\ell_k/k$. This implies $y_1(k) = \ell_k R_1(k)$ and $y_2(k) = \ell_k R_2(k)$ where

$$R_1(k) = \frac{\delta[\delta L(\omega_k/\delta) - L(\omega_k)]}{\omega_k^2}, \tag{49}$$

$$R_2(k) = \frac{\delta[L(\omega_k/\delta) - L(\omega_k)]}{\omega_k}. \tag{50}$$

Since $\omega_k \to 0$ as $k \to \infty$, l'Hôpital's rule yields

$$\lim_{k\to\infty} R_1(k) = \lim_{z\to 0} \frac{\delta[\delta L(z/\delta) - L(z)]}{z^2} = -\frac{1-\delta}{2}, \tag{51}$$

$$\lim_{k\to\infty} R_2(k) = \lim_{z\to\infty} \frac{\delta[L(z/\delta) - L(z)]}{z} = 1-\delta. \tag{52}$$

It follows from (54) that

$$P^{(k)}(s^*) = -\frac{\ell_k}{1-\delta}[R_1(k) + R_2(k)]. \tag{53}$$

Moreover, (51) and (52) imply that for any $\eta_0 > 0$, there exists K_0 such that

$$P^{(k)}(s^*) = -\frac{\ell_k}{1-\delta}\left(\frac{1-\delta}{2} - \eta_0\right), \quad k \ge K_0. \tag{54}$$

The claim then follows. ■