

AGE OF INFORMATION WITH FINITE HORIZON AND PARTIAL UPDATES

David Ramirez^{*}

Elza Erkip[†]

H. Vincent Poor^{*}

^{*}Department of Electrical Engineering, Princeton University, [†]New York University

ABSTRACT

A resource-constrained system monitors a source of information by requesting a finite number of updates subject to random transmission delays. An a priori fixed update request policy is shown to minimize a polynomial penalty function of the age of information over arbitrary time horizons. Partial updates, compressed updates with reduced transmission and information content, in the presented model are shown to incur an age penalty independent of the compression. Finite horizons are shown to have better performance in terms of second order statistic relative to infinite horizons.

Index Terms— Age of information, remote sensing, wireless communication

1. INTRODUCTION

Tactile internet is on the horizon for wireless network services, if highly responsive connectivity between environment-facing sensors and user-facing systems is available [1]. Sensor updates must be sparse enough to extend battery life yet frequent enough to satisfy a given system performance requirement. A pressing question, and the focus of this work, is to find an update request policy that gracefully maintains system performance of a resource constrained system.

When actions depend on the most recently sensed data, system performance degrades as the elapsed time since the last received update grows. The concept of age of information (AoI), as presented by Kaul et al. [2], quantifies the system performance penalty associated with the delay between received updates. AoI is a meaningful metric which differs from delay and throughput by focusing on the time between update generation and update arrival [3], i.e., AoI is a receiver-centric measure of delay. Achieving minimal AoI can be nontrivial, e.g., [4] showed constraining the number of updates makes a zero-wait policy sub-optimal in terms of AoI.

Results in [5] characterized the impact of various queuing strategies on the peak AoI. In [6] it was shown that an AoI threshold strategy minimizes the average AoI under long-term average update constraints and error-correction techniques. Work in [7] showed that with adequate queuing strategies the impact of buffering capacity of intermediary relays becomes nearly negligible. Non-linear AoI penalties for a randomly

update generating source were considered in [8]. A coding-inspired approach where updates vary in information quantity and bit size was presented in [9]. A system monitoring a random source with a finite number of updates was considered in [10] where a threshold based strategy optimally minimized the expected estimation error. A controllable update processing time was considered in [11] to minimize AoI constrained to a distortion metric inversely proportional to the processing time of updates sent over a deterministic channel. Mutual information (MI) of a source and destination was proposed as a measure for AoI in [12], where an MI thresholding strategy was shown to be optimal. Recently, [13] considered a central scheduler tasked with minimizing the total AoI by controlling how updates are relayed from sources to destinations.

The present work is distinguished via the use of a model that allows for new update requests to be made even if a previous update has not yet arrived, operates over arbitrary time horizons, and considers the use of partial updates. The AoI penalty is a polynomial function of arbitrary degree and finite first and second moments are assumed computable for the random delay. Minimization of the total AoI penalty is obtained by a priori computing update request times. The impact on AoI penalty of reordering update requests is presented. Additionally, it is shown that a finite horizon approach achieves smaller variance of AoI penalty, thus greater stability can be provided to systems with high reliability requirements.

2. NETWORK MODEL

A monitoring system requests information updates via wireless links from a source of information (e.g., a set of sensors). The system is limited to making at most $N \in \mathbb{Z}^+$ update requests over a time period $T \in \mathbb{R}^+$. Note that $\frac{N}{T}$ can serve as an average update constraint or a constraint on energy if an energy per update ratio is introduced.

Define $i \in \{1, \dots, N\}$ as a labeling of the N updates, $t : 0 \leq t \leq T, t \in \mathbb{R}$ as a time index, $\delta_i \in [0, T]$ as the departure time of the i -th update request. *W.l.o.g.*, labels are such that $\delta_i \leq \delta_j \forall i < j \in \{1, \dots, N\}$. Define d_i as the transmission time of update i , modeled as a continuous strictly positive random variable with finite first and second moments. Statistics of all transmission times are assumed independent from each other, but not identical to allow the model to account for the N updates to arrive from potentially N distinct sensors.

Work partially done while Ramirez was at New York University.

Define $\alpha_i = \delta_i + d_i$ as the arrival time of update i .

To account for, either, multiple paths between system and sensors or higher layer mechanisms giving sequential requests different delays (e.g., exponential backoff in 802.11), assume new update requests can be made even if a previous update has not yet arrived. At time t , define the *freshes* update $u(t)$ as $\arg \max_i \delta_i : \alpha_i \leq t$ with $u(t) = 0 \forall t < \alpha_1$, i.e., the freshest update is a received update that was generated most recently. Update i is said to be useful if $\exists t : i = u(t)$ else the update is wasted.

Define $\delta = \{\delta_i \forall i\}$ as a policy, and note that the system only controls what policy is used. Policy selection impacts the system penalty function, defined as a function which is discontinuous at reception of useful updates and otherwise monotonically increasing with time. Related literature names this penalty function the *age of information* and is defined as

$$A_k(t, \delta) = C_k(t - \alpha_{u(t)} + d_{u(t)})^k \quad (1)$$

with $k \geq 1$, $C_k \in \mathbb{R}^+$ as a scaling value, $A_0 \in \mathbb{R}^+$ as an initial age, and $d_0 = \left(\frac{A_0}{C_k}\right)^{1/k}$. Define the set of useful updates $\alpha^u = \{\alpha_j \forall j : \exists t, u(t) = j\} \cup \{a_0^u = 0, a_{N+1}^u = T\}$ and the total penalty can be defined as

$$\mathcal{S}_k(\delta) = \sum_{i=0}^N \int_{\alpha_i^u}^{\alpha_{i+1}^u} A_k(t, \delta) dt, \quad (2)$$

where $A_k(t, \delta)$ is evaluated from inside the integration period. Note that $\mathcal{S}_k(\delta)$ is a random variable.

If a policy δ achieves $\mathcal{S}_k(\delta)$ with N updates, then there exists a policy with $N + 1$ updates with no larger a total penalty (e.g., the same N departure times and random departure time of the extra update). Furthermore, if a policy with $N + 1$ *useful* updates exists, then a lower total penalty can be guaranteed. The problem of interest is to obtain an optimal policy, i.e.,

$$\delta^* = \arg \min E[\mathcal{S}_k(\delta)]. \quad (3)$$

Note, an optimal policy has a penalty no greater than any other policy when averaged over *many* instances of length T .

3. OPTIMAL SOLUTION: CRITICAL AGE POLICY

First, a set of feasible solutions for (3) for a given limited expected peak age is defined. Then, a proof for the existence of an optimal solution and method to find an optimal policy are presented.

To account for discontinuities due to update reception in (1), define $f^-(t) = \lim_{t' \rightarrow t^-} f(t')$ and $f^+(t) = \lim_{t' \rightarrow t^+} f(t')$ for some function $f(t)$. Define an arbitrary limit on the expected peak age as the critical age A^C . For a critical age $A^C > A_k^-(E[d_1], \delta)$, the critical policy is defined as

$$\delta^C = \{\delta_i^C : E[A_k^-(t, \delta^C)] \leq A^C \forall t\}, \quad (4)$$

else if $A^C \leq A_k^-(E[d_1], \delta)$, the critical policy δ^C is

$$\delta^C = \{\delta_1^C = 0, \delta_i^C : E[A_k^-(t, \delta^C)] \leq A^C \forall t > \alpha_1\}, \quad (5)$$

where (5) sets $\delta_1^C = 0$ to mitigate A_0 as quickly as possible since A_k is smaller than the age at the earliest expected arrival of the first update (i.e., if $\delta_1 = 0$ then $E[\alpha_1] = E[d_1]$).

For a given N , define the set of all critical policies as \mathcal{P}_N . A critical age of $0 < A^C < \max_i E[d_i]$ cannot be guaranteed by any critical policy. Thus, among all policies in \mathcal{P}_N there exists a minimal A^C . Define $\delta^* := \arg \min_{\delta \in \mathcal{P}_N} A^C$ as the critical policy with smallest critical age $A^* := \min_{\delta \in \mathcal{P}_N} A^C$. Theorem 1 asserts that a critical policy solves (3).

Theorem 1. *A critical policy $\delta^* \in \mathcal{P}_N$ obtains an average total penalty no greater than any other policy.*

Proof. Define a policy that differs from δ^* in at least one element, i.e., $\delta' : \exists \delta'_i \neq \delta_i^*$. Furthermore, define a subset of all the differing elements as $M = \{m : \delta'_m \neq \delta_m^*, 1 \leq m \leq N\}$. Assume that δ' and δ^* have equal numbers of useful updates, else δ' can be improved upon. Define $\Delta_m = (\delta'_m - \delta'_{m-1}) - (\delta_m^* - \delta_{m-1}^*) \forall m \in M$, with $\delta'_0 = \delta_0^* = 0$ for completeness, as the temporal shift between the policies for consecutive departures of updates in the set M .

Delay, i.e. $d_i \forall i$, is a policy independent stochastic process, thus $A_k^+(\alpha_m^*, \delta^*) = A_k^+(\alpha'_m, \delta')$ and the penalty incurred by δ' for $t \in [\alpha'_m, \min(\alpha_{m+1}^* - \Delta_m, \alpha'_{m+1})]$ is equally incurred by policy δ^* for $t \in [\alpha_m^*, \min(\alpha_{m+1}^*, \alpha'_{m+1} - \Delta_m)]$.

Focus on the penalty incurred during the remaining time period of length $|\Delta_m|$. First, if $|\Delta_m| = 0$ then the impact of update m for policy δ' is equivalent and time shifted version of the impact of update m for policy δ^* . Note that $\exists m : |\Delta_m| \neq 0$ since both policies have an equal and finite number of useful updates and T is finite. Note that the time period of length $|\Delta_m|$ need not be a continuous subset of the time period $[0, T]$.

Second, if $\delta_1^* \neq 0$, then over *any* time period of length $|\Delta_m|$ the policy δ^* incurs a total penalty no greater than $|\Delta_m|A^C$. Alternatively, if $\delta_1^* = 0$ the total penalty incurred by both policies over time period $[0, \alpha_1]$ is equivalent, and for δ^* no greater than $|\Delta_m|A^C$ for any period of length $|\Delta_m|$ in the interval $[\alpha_1, T]$.

Third, for any $\Delta_m > 0$ (i.e., transmitting later than δ^*) policy δ' incurs a total penalty no less than $|\Delta_m|A^C$ at a time period after α_m^* since the penalty function is monotonically increasing. For any $\Delta_m < 0$ (i.e., transmitting earlier than δ^*) policy δ' incurs a penalty after α'_m equivalent to the penalty incurred by δ^* after α_m^* ; yet since T and N are finite there exists a period of length $|\Delta_m|$ for which δ' incurs a penalty greater than $|\Delta_m|A^C$. \square

Essentially, any deviation from the optimal policy δ^* creates a temporal shift of the penalty function. Since the penalty function is non-negative and positive monotonic, a




Fig. 1: Example of possible differences between δ^* and δ' . Updates m and $m + 1$ cause a greater penalty in δ' , while update $m + 2$ does not by itself cause a penalty difference.

larger penalty is incurred. Fig. 1 illustrates possible values of Δ_m as described in the previous proof.

Construction of δ^* provides one equation for each of the N departure times δ_i with $N + 1$ variables, i.e., the N departure times plus the critical age A^C . The additional equation necessary to make the system solvable comes from realizing that at time T the age should also not exceed A^C . Thus a critical policy δ^* satisfies the following:

$$A_k^-(\delta_i^* + E[d_i], \delta^*) = A^* \forall i \text{ and } A_k^-(T, \delta^*) = A^*, \quad (6)$$

from which one can obtain the value of

$$A^* = C_k \left(\frac{A_0 + T + \sum_{i=1}^N E[d_i]}{N + 1} \right)^k. \quad (7)$$

Furthermore, one can obtain

$$\delta_N^* = \frac{NT - A_0 - \sum_{i=1}^N E[d_i]}{N + 1} \quad (8)$$

and then recursively solve for the remaining $N - 1$ values. Finally, the average total penalty incurred by policy δ^* is

$$\overline{S}_k(\delta^*) = \sum_{i=0}^N \frac{C_k \left[(\delta_{i+1}^* + E[d_{i+1}] - \delta_i^*)^{k+1} - E[d_i]^{k+1} \right]}{(k + 1)}, \quad (9)$$

with $\delta_0^* = d_0 = -A_0^{1/k}$, $d_{N+1} = 0$, and $\delta_{N+1}^* = T$ defined for completeness and compactness.

For large values of A_0 or $E[d_i]$, e.g., $A_0 > NT - \sum_{i=1}^N E[d_i]$, optimal departure times may become negative, cf. (8). In such cases, the first update is $\delta_1 = 0$ and the problem is recast with the remaining $N - 1$ departures. While optimal departure times for any value of k are equivalent, the critical age of the policy is distinct.

Augmenting a policy with additional *useful* updates decreases the total age penalty. Yet, reordering update requests

can turn a useful update into a wasteful update. In the simple case when all N update requests experience statistically equivalent delays, the time between any two consecutive departures is constant. Formally, if $E[d_i] = E[d_j] \forall i, j$ then $\delta_i^* - \delta_{i-1}^* = \delta_j^* - \delta_{j-1}^* \forall i, j$. Meaning, if all updates are transmitted over the same stationary wireless channel, then a constant time period between update requests is optimal and reordering updates does not impact the solution of (3).

Now, consider only the case when $\exists i, j : E[d_i] \neq E[d_j]$. First, if $k = 1$ and $A_0 = 0$ then any achievable value of A^C by a policy δ^* is achievable by any reordering of the update requests with a different policy. This follows from (7) being independent of δ^* and $A_0 = 0$. Second, if $A_0 \gg 0$ the first departure may be forcefully fixed to $\delta_1 = 0$ to mitigate the penalty as soon as possible. Intuitively, when reordering is possible and a value of A_0 is large enough that $\delta_1^* = 0$, then the first update should be requested from the source with minimum delay, i.e., reorder such that $\arg \min_i E[d_i] = 1$. Third, if only the first element is forcefully fixed due to a large initial penalty, then the remaining $N - 1$ elements can achieve equivalent average total age penalty with any reordering. Reordering, as described, does not impact A^* but does impact the values of $\delta_i \in \delta^*$.

4. PARTIAL UPDATES

In this section, and only this section, consider that an update can be compressed to reduce transmission time and information content (i.e., lossy compression). The compressed update is a partial update, i.e., for update i with transmission time d_i has a corresponding partial update with transmission time d_p such that $E[d_p] < E[d_i]$ and $A_k^+(\alpha_i, \delta) < A_k^+(\alpha_p, \delta)$ if $A_k^-(\alpha_i, \delta) = A_k^-(\alpha_p, \delta)$. Compare with [9] where a partial update was a lossless compressed version of a full update.

Assume a partial update *resets* the AoI to the AoI at reception of the previous useful update. Meaning, for a partial update $p \in \{1, \dots, N\}$ at time t such that $p = u(t)$ define the previous update arrival time as $t' = \lim_{\epsilon \rightarrow 0} \alpha_{u(t-\epsilon)}$ with $\alpha_0 = 0$, and $A_k^+(\alpha_p, \delta) := A_k^+(t', \delta)$. Thus, the partial update provides no *fresher* information than that of the previous update. Such a model for partial update inherently favors the use of partial updates over a full update whenever the expected delay increases the AoI penalty no greater than the penalty at the previous update arrival.

An optimal critical policy (cf. Section III) using only partial updates, i.e., δ_p , can lead to a total penalty that is independent of transmission times. E.g., with $k = 1$ the optimal critical policy creates a sawtooth wave of $N + 1$ peaks that go from A_0 up to $A_0 + C_k \frac{T}{N+1}$ thus obtaining a total penalty of

$$S_k(\delta_p) = \frac{C_k T^2}{2(N + 1)} + A_0 T. \quad (10)$$

Note the lack of dependence of (10) on the update delay, meaning that reduced delay of a partial update does not im-




Fig. 2: Avg. Penalty per Horizon $S_k(\delta^*)/T$ as a function of T . For $T > 100$, the Unknown Horizon outperforms the Infinite Horizon, and at $T > 250$, the scenarios are comparable.

part the performance. When the delay of full updates is $d_i \leq \frac{A_0}{C} \forall i$, then a policy of only full updates will not incur a larger total linear penalty than a policy of only partial updates.

5. IMPACT ON SECOND ORDER STATISTICS

Fixing N and increasing T leads to an increase in (3), while fixing T and increasing N can decrease (3). Fixing the ratio $\frac{N}{T}$ allows for arbitrary, yet equal, growth in N and T which maintains equivalent average resource consumption. From (7) one can infer that increasing N and T at the same rate will drive A^* towards $C_k \left(\frac{T}{N} + \overline{E[d_i]} \right)^{1/k}$ where $\overline{E[d_i]}$ is the arithmetic mean of expected delays.

From (9), note that increasing N and T by equal amounts leads to an increase in the total penalty $S_k(\delta^*)$. Adding to the growth in average penalty is the negative impact on the variance of the penalty experienced over many realizations, defined as $\sigma_{S_k(\delta^*)}^2$. With (9) upper and lower bounded by TA^* and $\min(A_0, E[d_i] \forall i)$, respectively, one can use Bhatia-Davis' inequality [14] to upper bound the variance as

$$\sigma_{S_k(\delta^*)}^2 < (TA^* - \overline{S_k(\delta^*)}) (\overline{S_k(\delta^*)} - \min(A_0, E[d_i] \forall i)) \quad (11)$$

where various refinements can be done if additional assumptions are made [15]. Now, if increasing N and T is done such that the additional update requests do not reduce the value of $\min(A_0, E[d_i] \forall i)$, then one can expect the variance of the total penalty incurred to increase with N and T . Therefore, a shorter horizon (i.e., smaller T) can lead to smaller values of σ^2 , and ultimately more stable performance from the system.

The elapsed time before the system is set to *initial conditions* defines the horizon T . Defining *initial conditions* is where system engineering comes into play. A system where delay is temporally stationary has a horizon $T = \infty$, while a system where the estimated statistics on delay are applicable for a finite amount of time before they need to be re-estimated

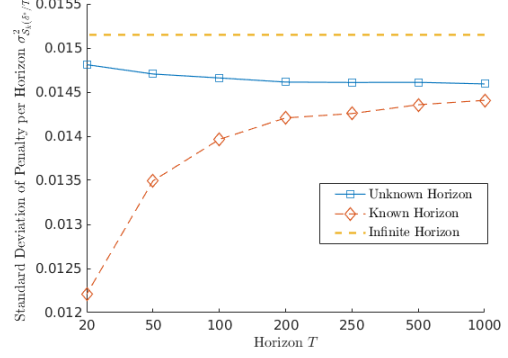


Fig. 3: Standard deviation of $\sigma_{S_k(\delta^*)}^2$ as a function of T . Infinite Horizon underperforms among evaluated scenarios.

has a horizon $T < \infty$. Horizon selection impacts second order statistics, as suggested by (11) and analyzed in the sequel.

6. SIMULATION RESULTS

Simulation parameters are set to $\frac{N}{T} = 0.4$, $A_0 = 0$, $k = 1$, $C_k = 1$, and $d_i \sim U(0, 1)$. Three scenarios are evaluated. Infinite Horizon solves (3) with $T = \infty$. Known Horizon uses T as the horizon. Unknown Horizon concatenates periods of length $\frac{1000}{T}$, using $\frac{N}{T}$ updates per period, and the penalty at the end of one period is the initial age of the next period.

Fig. 2 shows $E[\frac{S_k(\delta^*)}{T}]$ as a function of T , where a lower value on the vertical axis implies better performance. Intuitively, using larger T moves all scenarios towards the Infinite Horizon scenario. Note the existence of a trade-off point ($T < 100$) where assuming an infinite horizon is a better strategy than estimating an unknown horizon. The Unknown Horizon scenario suffers from the initial age carried over from previous periods, thus its performance improves with increase in T . When the horizon is known, a smaller horizon allows for a finer control of resources, thus a smaller penalty. In Fig. 3 the standard deviation of $\frac{S_k(\delta^*)}{T}$ is shown as a function of T , where a lower value implies better performance. Assuming an infinite horizon leads to a larger standard deviation, and a greater variation in performance.

7. CONCLUSION

An update request policy that minimizes the average information age penalty of a resource constrained monitoring system with stochastic update delay has been presented. The optimal departure times have been obtained by solving a linear system of equations which grows with the available number of updates. Results have also outlined that using partial updates incurs a total age independent of delay, but is only beneficial when the delay reduction outweighs the increase in age.

8. REFERENCES

- [1] M. Simsek, A. Aijaz, M. Dohler, J. Sachs, and G. Fettweis, "5G-Enabled Tactile Internet," *IEEE J. on Sel. Areas in Commun.*, vol. 34, no. 3, pp. 460–473, 2016.
- [2] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing Age of Information in Vehicular Networks," in *Proc. Conf. on Sensor, Mesh and Ad Hoc Commun. and Netw.* IEEE, 2011.
- [3] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Minimizing the Age of Information in Broadcast Wireless Networks," in *IEEE Proc. 54th Annu. Allerton Conf. on Commun., Control, and Computing*, 2016.
- [4] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or Wait: How to Keep Your Data Fresh," *IEEE Trans. Inf. Theory*, vol. 63, no. 11, pp. 7492–7508, 2017.
- [5] K. Chen and L. Huang, "Age-of-Information in the Presence of Error," in *IEEE Int. Symp. on Inf. Theory*, 2016, pp. 2579–2583.
- [6] E. T. Ceran, D. Gunduz, and A. Gyorgy, "Average Age of Information with Hybrid ARQ under a Resource Constraint," *arXiv preprint arXiv:1710.04971*, 2017.
- [7] Y.-P. Hsu, E. Modiano, and L. Duan, "Age of Information: Design and Analysis of Optimal Scheduling Algorithms," in *IEEE Int. Symp. on Inf. Theory*, 2017, pp. 561–565.
- [8] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "Age and Value of Information: Non-linear Age Case," in *IEEE Int. Symp. on Inf. Theory*, 2017, pp. 326–330.
- [9] S. Poojary, S. Bhambay, and P. Parag, "Real-time Status Updates for Correlated Source," in *Inf. Theory Workshop*. IEEE, 2017.
- [10] J. Chakravorty and A. Mahajan, "Distortion-transmission trade-off in real-time transmission of gauss-markov sources," in *Intl. Symp. on Inf. Theory*. IEEE, 2015, pp. 1387–1391.
- [11] M. Bastopcu and S. Ulukus, "Age of information for updates with distortion," *arXiv preprint arXiv:1904.10444*, 2019.
- [12] Y. Sun and B. Cyr, "Information aging through queues: A mutual information perspective," in *Intl. Workshop on Signal Process. Advances in Wireless Commun.* IEEE, 2018, pp. 1–5.
- [13] O. Ayan, M. Vilgelm, M. Klügel, S. Hirche, and W. Kellerer, "Age-of-information vs. value-of-information scheduling for cellular networked control systems," in *Proc. ACM/IEEE Intl. Conf. on Cyber-Physical Syst.*, 2019, pp. 109–117.
- [14] R. Bhatia and C. Davis, "A Better Bound on the Variance," *The Amer. Math. Monthly*, vol. 107, no. 4, 2000.
- [15] R. Agarwal, N. S. Barnett, P. Cerone, and S. S. Dragomir, "A survey on some inequalities for expectation and variance," *Comput. & Math. with Appl.*, vol. 49, no. 2-3, pp. 429–480, 2005.