Who Should Google Scholar Update More Often?

Melih Bastopcu Sennur Ulukus
Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742
bastopcu@umd.edu *ulukus@umd.edu*

***Abstract*—We consider a resource-constrained updater, such as Google Scholar, which wishes to update the citation records of a group of researchers, who have different mean citation rates (and optionally, different importance coefficients), in such a way to keep the overall citation index as up to date as possible. The updater is resource-constrained and cannot update citations of all researchers all the time. In particular, it is subject to a total update rate constraint that it needs to distribute among individual researchers. We use a metric similar to the age of information: the long-term average difference between the actual citation numbers and the citation numbers according to the latest updates. We show that, in order to minimize this difference metric, the updater should allocate its total update capacity to researchers proportional to the *square roots* of their mean citation rates. That is, more prolific researchers should be updated more often, but there are diminishing returns due to the concavity of the square root function. More generally, our paper addresses the problem of optimal operation of a resource-constrained sampler that wishes to track multiple independent counting processes in a way that is as up to date as possible.**

I. Introduction

Consider a citation index such as Google Scholar. As abstracted out in Fig. 1, Google crawls the web to find and index various items such as documents, images, videos, etc. Focusing on scientific documents, Google Scholar further examines the contents of these documents to extract out citation counts for indexed papers. Google Scholar then needs to update citation counts of individual researchers, which there are many. We model the citation count of each individual researcher as a counting process with a fixed mean, e.g., λ_i for researcher i. Assuming that Google Scholar is resource-constrained, i.e., that it cannot update all researchers all the time, how should it prioritize updating researchers? If it can update only a fraction of all researchers, who should it update? Should it update researchers with higher mean citation rates more often as their citation counts are subject to larger change per unit time? Or should it update researchers with lower mean citation rates more often in order to capture rarer more informative changes?

We view this problem with the lens of the recent literature on *information freshness* quantified through the metric of age of information. Freshness, and age, of information have been studied in the context of web crawling [1]–[4], social networks [5], queueing networks [6]–[15], caching systems [16], [17], remote estimation [18]–[21], energy harvesting systems [22]–[31], fading wireless channels [32], [33], scheduling in networks [34]–[41], multi-hop multicast networks [42]–[45], lossless and lossy source coding [46]–[50], computation-intensive systems [51]–[56], vehicular, IoT and UAV systems [57], [58], reinforcement learning [59]–[61] and so on.

This work was supported by NSF Grants CCF 17-13977 and ECCS 18-07348.

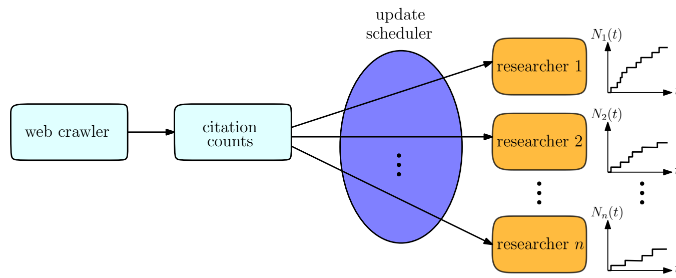

Fig. 1. Web crawler finds and indexes scientific documents, from which citation counts are extracted upon examining their contents. Scheduler schedules updating citation counts of individual researchers based on their mean citations, and optionally, importance factors, subject to a total update rate.

We model the citation count of each researcher as a counting process with a given mean. In particular, we model citation arrivals for researcher i as a Poisson counting process with rate λ_i. Optionally, we may further assign an importance factor to each researcher, based on their research field or citation count, but this is optional, and does not affect the structure of the results. If an importance coefficient is used, we denote it with μ_i for researcher i. Ideally the updater should update all researchers all the time, however, due to computational limitations, this may not be possible. We model the updater as a resource-constrained entity which has a total update capacity of c, which it should distribute among all researchers. We allocate an update rate ρ_i for updating researcher i. These ρ_i are collectively subject to the total system update capacity of c. We consider the cases of Poisson updates (i.e., updates with exponential inter-update times), deterministic updates, and synchronized updates. We determine the optimal update rates ρ_i subject to the total update rate c in a way to maximize the system freshness. In a broader sense, this problem is related to the problem of real-time timely estimation of signals, which have different change rates and importance factors, with the goal of finding the optimal individual sampling rates, under a total system sampling rate constraint. In this paper, we specialize this broader goal to the setting of counting processes and to the context of tracking citation counts of researchers.

References that are most closely related to our work are [4] and [18]. Reference [4] considers the problem of finding optimal crawl rates to keep the information in a search engine fresh while maintaining the constraints on crawling rates imposed by the websites and also the total crawl rate constraint of the search engine, in the presence of non-uniform importance scores and change rates for the websites. Reference [18] focuses on remote real-time reconstruction of a single Poisson counting process using uniform sampling. While taking more samples helps reconstruct the signal better, this increases queuing delays, which inherently affects the real-time signal estimation negatively. Reference [18] studies this trade-off and finds the optimal sampling rate. Our timeliness metric is similar to the one considered in [18], which is the difference of a counting process and its sampled (updated, in our case) version. However, we consider multiple counting processes with different arrival rates and importance factors, and optimize our update rates for all processes jointly under a total update rate constraint. Similar to [4], we consider exponential arrival and sampling times in a constrained manner, and allow for importance coefficients, however, our timeliness metric and our overall problem setting are different.

In this paper, we first find an analytical expression for the long-term average difference between the actual and updated counting processes. We then minimize this expression as a function of the update rates of individual researchers subject to the overall update rate. We show that the optimal update rates are proportional to the *square roots* of mean citation rates of the researchers for constant importance factors. Thus, it is optimal to update more prolific researchers more often, however, the proportionality is sub-linear and in the form of square root of the mean citation rate, i.e., there are diminishing returns due to the concavity (sub-linearity) of the square root function. We show that if the importance factors of the researchers are linear in their mean citation rates, then the optimal update rates are linear in their mean citation rates as well. We finally remark that other *square root* results have appeared in completely different settings in caching problems. In particular, it was found in [16] that, in order to keep the files fresh in a caching system, the (uniform) update rates of the files should be chosen proportional to the square roots of their popularity indices. In addition, it was shown in [62] that, in order to minimize the average number of websites searched, the number of websites that a file is mirrored should be chosen proportional to the square root of the file's popularity.

II. System Model and Problem Formulation

Consider n researchers. Let $N_i(t)$ denote the number of citations of researcher i. We model $N_i(t)$ as a Poisson process with rate λ_i, and assume that $N_i(t)$, for $i=1,\dots,n$, are independent. Let $t_{i,j}$ denote the time instance when the updater updates the number of citations of researcher i for the jth time. We denote the inter-update time between the jth and $(j-1)$th updates for researcher i as $\tau_{i,j}$. Based on these samples, the updater generates a real-time estimate of the counting process $N_i(t)$ as $\hat{N}_i(t)$, where

$$\hat{N}_i(t) = N_i(t_{i,j-1}), \quad t_{i,j-1} \leq t < t_{i,j}. \tag{1}$$

Fig. 2 shows $N_i(t)$ and $\hat{N}_i(t)$ with black and blue lines.

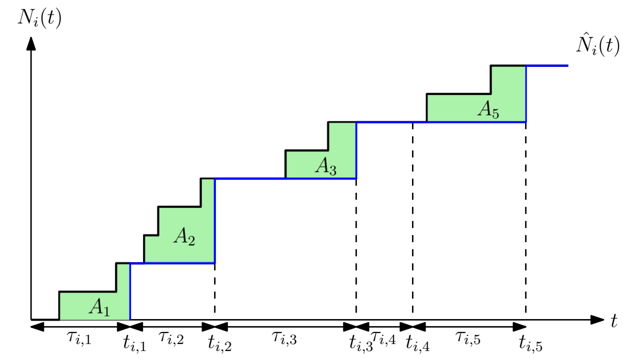

Fig. 2. The number of citations, $N_i(t)$, and the estimated number of citations, $\hat{N}_i(t)$, for researcher i. A_j denotes the total estimation error in $[t_{i,j-1}, t_{i,j})$.

Similar to [18], we use the average difference between the actual and updated processes as a measure of timeliness,

$$\Delta_i(T) = \frac{1}{T}\int_0^T \left(N_i(t) - \hat{N}_i(t)\right) dt. \tag{2}$$

If there are m updates in the interval $[0,T]$, then $\Delta_i(T)$ is

$$\Delta_i(T) = \frac{1}{T}\sum_{j=1}^{m} A_j, \tag{3}$$

where A_j is the difference in the interval $[t_{i,j-1}, t_{i,j})$, see Fig. 2. Then, the long term average difference for researcher i is $\Delta_i = \lim_{T\to\infty}\Delta_i(T)$, and can be written as [13],

$$\Delta_i = \lim_{T\to\infty}\frac{1}{T}\sum_{j=1}^{m} A_j = \lim_{T\to\infty}\frac{m}{T}\cdot\frac{1}{m}\sum_{j=1}^{m} A_j = \rho_i \mathbb{E}[A]. \tag{4}$$

Similar to the derivation in [18], conditioned on an arbitrary jth inter-update time, i.e., $\tau_i = \tau_{i,j} = d$, and the number of citation arrivals in that time interval $[t_{i,j-1}, t_{i,j-1}+d)$, i.e., $\tilde{N}_i(d) = N_i(t_{i,j-1}+d) - N_i(t_{i,j-1}) = k$, the expected difference is

$$\mathbb{E}\left[A|\tilde{N}_i(d) = k, \tau_i = d\right] = \frac{kd}{2}. \tag{5}$$

Thus, we have

$$\mathbb{E}\left[A|\tau_i = d\right] = \mathbb{E}\left[\mathbb{E}\left[A|\tilde{N}_i(d) = k, \tau_i = d\right]\right] = \frac{\lambda_i d^2}{2}. \tag{6}$$

In the following subsections, we present three different models for updating the citation numbers.

A. Model 1: Poisson Updater

In this model, shown in Fig. 3, the inter-update times for researcher i are exponential with rate ρ_i. Update processes for different researchers are independent. Continuing from (6), we find $\mathbb{E}[A]$ using exponential distribution as,

$$\mathbb{E}[A] = \int_0^\infty \mathbb{E}\left[A|\tau_i = t\right] f_{\tau_i}(t)dt = \frac{\lambda_i}{2}\mathbb{E}\left[\tau_i^2\right] = \frac{\lambda_i}{\rho_i^2}. \tag{7}$$

Thus, the long term average difference Δ_i in (4) with a Poisson updater is

$$\Delta_i = \frac{\lambda_i}{\rho_i}. \tag{8}$$

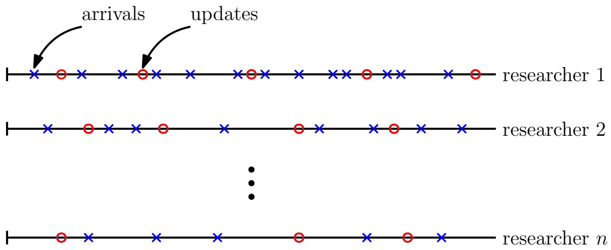

Fig. 3. Poisson updater: Inter-update times are exponential with rate ρ_i.

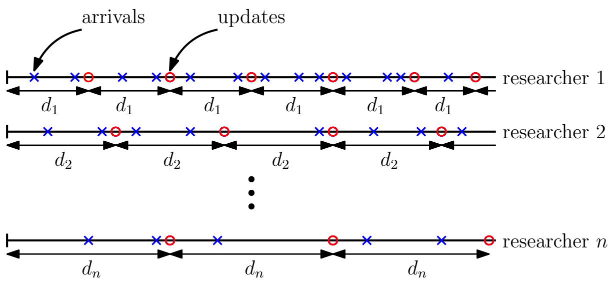

Fig. 4. Deterministic updater: Inter-update times are equal with $d_i = \frac{1}{\rho_i}$.

B. Model 2: Deterministic Updater

In this model, shown in Fig. 4, the inter-update times are deterministic and chosen optimally. Similar to [16], given that there are m_i updates for researcher i in the time interval $[0, T]$, the optimal inter-update times should be chosen equal to each other, i.e., $\tau_{i,j} = \frac{T}{m_i+1}$, for all j. Letting $T \to \infty$, this update scheme results in uniform sampling with rate ρ_i for researcher i where $\rho_i = \lim_{T\to\infty} \frac{m_i}{T}$. By using $d_i = \frac{1}{\rho_i}$ and (6), we obtain $\mathbb{E}[A] = \frac{\lambda_i}{2\rho_i^2}$. Thus, the long term average difference Δ_i in (4) with a deterministic (and uniform) updater is

$$\Delta_i = \frac{\lambda_i}{2\rho_i}. \tag{9}$$

C. Model 3: Common Synchronized Probabilistic Updater

In this model, shown in Fig. 5, the updater has a common synchronized update schedule that applies to all researchers. The inter-update times of the common updater are exponential with rate ρ. At each update instant, researcher i is updated with probability p_i independently of other researchers. Thus, inter-update times for researcher i are exponential with rate ρp_i. Note that here, we create the Poisson updates for researcher i by thinning the Poisson common updates using probabilistic updates according to p_i. The main problem, therefore, is to choose p_i for each researcher as it determines its mean update rate. This problem is the same as the one in Section II-A and in the optimal policy p_i is $p_i = \frac{\rho_i}{\rho}$ assuming ρ is sufficiently large to have feasible p_i, i.e., $0 \le p_i \le 1$, for all i.

D. Problem Formulation

Researcher i has the mean citation rate λ_i. In addition, and optionally, we consider an importance factor, μ_i, for researcher i. This may be removed by choosing all $\mu_i = \mu$. Then, the total long term average difference (over all researchers) becomes $\Delta = \sum_{i=1}^n \mu_i \Delta_i$, where per researcher difference, Δ_i, is

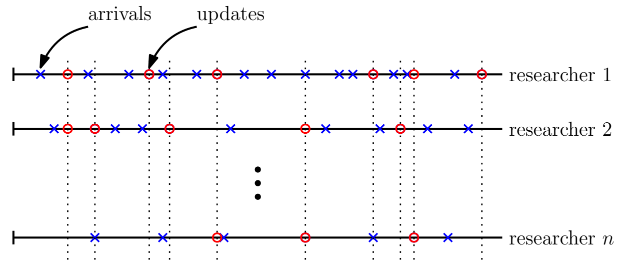

Fig. 5. Common synchronized updater: Common synchronized inter-update times are exponential with rate ρ. At each common update opportunity, researcher i is updated with probability p_i

given by (8) for the Poisson updater and common synchronized updater models, and by (9) for the deterministic updater model. The expressions in (8) and (9) differ only by a factor of 2, which is inconsequential for optimization purposes. Therefore, without loss of generality, from now on, we use the expression in (8). In addition, due to computational limitations, the updater is subject to a total update rate constraint $\sum_{i=1}^n \rho_i \le c$. Our aim is to find the optimal update rates for all researchers, ρ_i, for $i = 1, \dots, n$, such that the total long term average difference, Δ, is minimized while satisfying the constraint on the total update rate. Thus, our optimization problem is,

$$\begin{aligned} \min_{\{\rho_i\}} \quad & \sum_{i=1}^n \frac{\mu_i \lambda_i}{\rho_i} \\ \text{s.t.} \quad & \sum_{i=1}^n \rho_i \le c \\ & \rho_i \ge 0, \quad i = 1, \dots, n. \end{aligned} \tag{10}$$

We solve the optimization problem in (10) in the next section.

III. The Optimal Solution

The optimization problem in (10) is convex as the cost function is convex and the constraints are linear. We introduce the Lagrangian function [63] for (10) as

$$\mathcal{L} = \sum_{i=1}^n \frac{\mu_i \lambda_i}{\rho_i} + \beta \left(\sum_{i=1}^n \rho_i - c \right) - \sum_{i=1}^n \nu_i \rho_i, \tag{11}$$

where $\beta \ge 0$ and $\nu_i \ge 0$ for all i. Next, we write the KKT conditions as

$$\frac{\partial \mathcal{L}}{\partial \rho_i} = -\frac{\mu_i \lambda_i}{\rho_i^2} + \beta - \nu_i = 0, \tag{12}$$

for all i, and the complementary slackness conditions as

$$\beta \left(\sum_{i=1}^n \rho_i - c \right) = 0, \tag{13}$$

$$\nu_i \rho_i = 0, \tag{14}$$

for all i. Since the optimization problem in (10) is convex, the KKT conditions are necessary and sufficient.

First, we observe that the total update rate constraint $\sum_{i=1}^n \rho_i \le c$ must be satisfied with equality. If there is an update rate allocation policy such that $\sum_{i=1}^n \rho_i < c$, then we can achieve a lower average difference by increasing any ρ_i

as the cost function of (10) is a decreasing function of ρ_i. Thus, in the optimal update rate allocation policy, we must have $\sum_{i=1}^n \rho_i = c$ and $\beta \geq 0$ due to (13).

Next, we note that in the optimal policy, we must have $\rho_i > 0$, for all i, as $\rho_i = 0$ leads to infinite objective function in (10) which clearly cannot be an optimal solution. Thus, for the optimal rate allocations, we have $\rho_i > 0$ and $\nu_i = 0$, for all i, due to (14).

From (12), we find $\rho_i = \sqrt{\frac{\mu_i \lambda_i}{\beta}}$. By using $\sum_{i=1}^n \rho_i = c$, we solve $\beta = \frac{\left(\sum_{i=1}^n \sqrt{\mu_i \lambda_i}\right)^2}{c^2}$, which gives the optimal policy,

$$\rho_i = \frac{c\sqrt{\mu_i \lambda_i}}{\left(\sum_{j=1}^n \sqrt{\mu_j \lambda_j}\right)}, \quad i = 1, \dots, n. \tag{15}$$

Using the optimal rate allocation policy in (15), we obtain

$$\Delta_i = \frac{\sqrt{\lambda_i}\left(\sum_{j=1}^n \sqrt{\mu_j \lambda_j}\right)}{c\sqrt{\mu_i}}, \quad i = 1, \dots, n, \tag{16}$$

and the total long term average difference Δ as

$$\Delta = \frac{\left(\sum_{j=1}^n \sqrt{\mu_j \lambda_j}\right)^2}{c}. \tag{17}$$

Thus, the optimal update rates allocated to researchers in (15) are proportional to the square roots of their importance factors, μ_i, multiplied by their mean citation rates, λ_i. We note that if we ignore the importance factors, i.e., $\mu_i = \mu = 1$, then the optimal update rates are proportional to the square roots of the mean citation rates. On the other hand, if we choose the importance factors as proportional to the mean citation rates, i.e., $\mu_i = \alpha\lambda_i$, then the optimal update rates become linear in the mean citation rates.

IV. Numerical Results

In this section, we provide three numerical results. In the first two examples, we choose the mean citation rates as

$$\lambda_i = ar^i, \quad i = 1, \dots, n, \tag{18}$$

where $a > 0$ and $0 < r \leq 1$.

In the first example, we take $a = 10$, $r = 0.75$, $n = 20$ and $c = 10$. For this example, we use uniform importance coefficients, i.e., $\mu_i = 1$, for all i. We observe in Fig. 6(b) that researchers with higher mean citation rates have higher long term average difference Δ_i even though they are updated with higher update rates shown in Fig. 6(a). Further, we observe in Fig. 6(a) that due to diminishing returns caused by the square root allocation policy, update rates of the researchers with low mean citation rates are still comparable to the update rates of the researchers with high mean citation rates.

In the second example, we consider the case where the importance factors are chosen proportional to the mean citation rates of the researchers. We call such coefficients as λ-*proportional* importance coefficients, which are given by

$$\mu_i = \frac{\lambda_i}{\sum_{j=1}^n \lambda_j}, \quad i = 1, \dots, n. \tag{19}$$

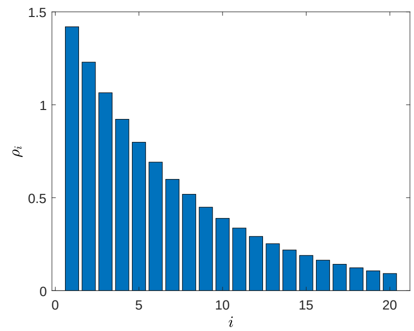

(a)

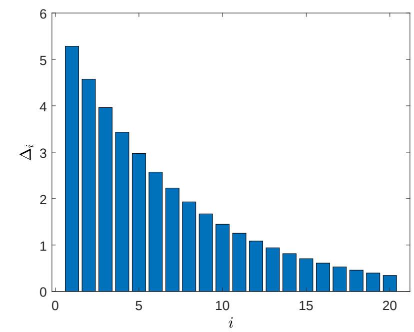

(b)

Fig. 6. (a) Optimal update rate allocation for each researcher, and (b) the corresponding optimal long term average difference Δ_i, when we use uniform importance coefficients $\mu_i = 1$, with λ_i given in (18), with $a = 10$ and $r = 0.75$ for $n = 20$.

In order to make a fair comparison between the λ-*proportional* and *uniform* importance coefficients, we scale the uniform importance coefficients as $\mu_i = \frac{1}{n}$, for all i.

As mentioned at the end of Section III, when we use λ-*proportional* importance coefficients, the optimal update rates become linear in the mean citation rates. We observe in Fig. 7(a) that using λ-*proportional* importance coefficients favor researchers with higher mean citation rates as their update rates increase compared to the update rates with the uniform importance coefficients. Further, we observe in Fig. 7(b) that the long term average differences are equal to each other when λ-*proportional* importance coefficients are used.

In the third example, we choose the mean citation rates as

$$\lambda_i = \frac{ar^i}{\sum_{j=1}^n r^j}, \quad i = 1, \dots, n, \tag{20}$$

which satisfy $\sum_{i=1}^n \lambda_i = a$. Note that, by this selection, we force total citation means of all researchers to be a constant. For this example, we take $n = 10$, $a = 1$ and consider three different r, which are $r = 0.5, 0.75, 1$. Note also that, a smaller r corresponds to a less even (more polarized) distribution of total mean citation rates among the researchers. We use uniform importance coefficients and plot achieved Δ with respect to c in Fig. 8. We observe in Fig. 8 that more polarized distribution of mean citation rates (smaller r) yields a lower Δ for the system, as we exploit the differences among the researchers by

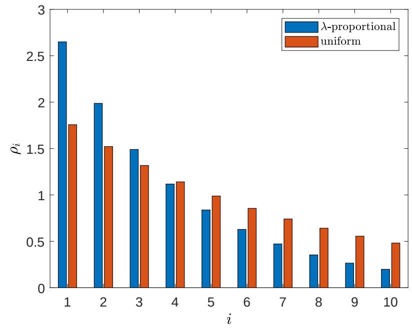

(a)

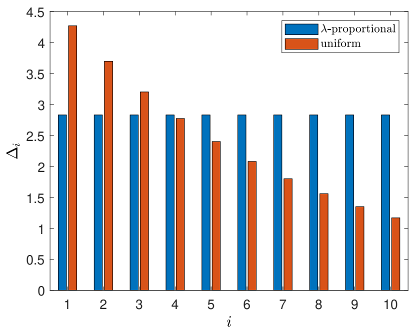

(b)

Fig. 7. (a) Optimal update rate allocation for each researcher, and (b) the corresponding optimal long term average difference Δ_i, when we use λ-*proportional* and uniform importance coefficients, with λ_i given in (18), with $a = 10$ and $r = 0.75$ for $n = 10$.

allocating even higher update rates to researchers with higher mean citation rates. As an aside, we note that if we used λ-*proportional* importance coefficients, we would have a Δ which is independent of individual λ_i that depends only on the sum of λ_i which is a here. This achieved Δ is also equal to the Δ achieved with uniform importance coefficients when $r = 1$ which is shown as the blue dashed line in Fig. 8. Thus, if we use λ-*proportional* importance coefficients, the achieved Δ is independent of the mean citation rate distribution among the researchers, but it results in higher Δ. In other words, uniform importance coefficients achieve lower Δ compared to λ-*proportional* importance coefficients in this case for $r < 1$.

V. Conclusion, Discussion and Future Directions

We considered the problem of timely updating of citation counts by a resource-constrained updater. We showed that the optimal policy is to choose the update rates of individual researchers proportional to the square roots of their mean citation rates multiplied by their importance factors (if any).

Next, we discuss limitations of our model and suggest future research directions. First, we note that we modeled the citation numbers of a researcher as a counting process, which is monotonically increasing and increments one at a time. The monotonically increasing nature of the counting process assumes that published articles are always available and cannot be withdrawn or changed, which may not be the case, especially if the crawler cannot reach certain websites that it used to reach, and lose access to previously counted publications. In addition, one at a time increments assume that citations come one by one, which may not be the case as conferences publish all the articles simultaneously in proceedings, and journals publish articles in monthly issues. That is, the increments in citations for each researcher may be more than one at a time. Thus, modeling the number of citations with a simple Poisson process may not be sufficient. However, as the researchers increasingly publish their works as they are completed on their personal websites or academic websites such as arXiv, it may still be acceptable to model arrivals as Poisson processes. Even then though, multiple citations to the same researcher from a single article will not be captured by the model in this paper. In addition, considering the fact that researchers often collaborate and publish articles jointly, the arrivals of citations for different researchers might be correlated. In this paper, we considered the case where the citation arrivals for researchers are independent counting processes. This may be feasible by focusing on researchers with no common publications.

Finally, we assumed that mean citation rates and importance factors are fixed and known by the updater. As a future direction, an online setting can be considered where both of these parameters are learned from observations over time.

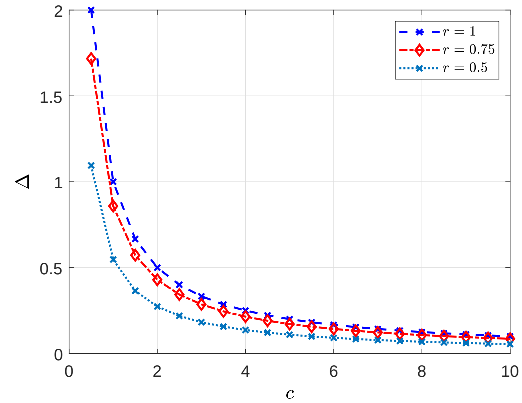

Fig. 8. Total long term average difference Δ with respect to c, when uniform importance coefficients are used and λ_i are given in (20), with $a = 1$ and $r = 0.5, 0.75, 1$ for $n = 10$.

References

[1] J. Cho and H. Garcia-Molina. Effective page refresh policies for web crawlers. *ACM Transactions on Database Systems*, 28(4):390–426, December 2003.

[2] B. E. Brewington and G. Cybenko. Keeping up with the changing web. *Computer*, 33(5):52–58, May 2000.

[3] Y. Azar, E. Horvitz, E. Lubetzky, Y. Peres, and D. Shahaf. Tractable near-optimal policies for crawling. *PNAS*, 115(32):8099–8103, August 2018.

[4] A. Kolobov, Y. Peres, E. Lubetzky, and E. Horvitz. Optimal freshness crawl under politeness constraints. In *ACM SIGIR Conference*, pages 495–504, July 2019.

[5] S. Ioannidis, A. Chaintreau, and L. Massoulie. Optimal and scalable distribution of content updates over a mobile social network. In *IEEE Infocom*, April 2009.

[6] S. K. Kaul, R. D. Yates, and M. Gruteser. Real-time status: How often should one update? In *IEEE Infocom*, March 2012.

[7] M. Costa, M. Codrenau, and A. Ephremides. Age of information with packet management. In *IEEE ISIT*, June 2014.
[8] A. M. Bedewy, Y. Sun, and N. B. Shroff. Optimizing data freshness, throughput, and delay in multi-server information-update systems. In *IEEE ISIT*, July 2016.
[9] Q. He, D. Yuan, and A. Ephremides. Optimizing freshness of information: On minimum age link scheduling in wireless systems. In *IEEE WiOpt*, May 2016.
[10] C. Kam, S. Kompella, G. D. Nguyen, Wieselthier J. E., and A. Ephremides. Age of information with a packet deadline. In *IEEE ISIT*, July 2016.
[11] 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 Transactions on Information Theory*, 63(11):7492–7508, November 2017.
[12] E. Najm and E. Telatar. Status updates in a multi-stream M/G/1/1 preemptive queue. In *IEEE Infocom*, April 2018.
[13] E. Najm, R. D. Yates, and E. Soljanin. Status updates through M/G/1/1 queues with HARQ. In *IEEE ISIT*, June 2017.
[14] A. Soysal and S. Ulukus. Age of information in G/G/1/1 systems. In *Asilomar Conference*, November 2019.
[15] A. Soysal and S. Ulukus. Age of information in G/G/1/1 systems: Age expressions, bounds, special cases, and optimization. Available on arXiv: 1905.13743.
[16] R. D. Yates, P. Ciblat, A. Yener, and M. Wigger. Age-optimal constrained cache updating. In *IEEE ISIT*, June 2017.
[17] H. Tang, P. Ciblat, J. Wang, M. Wigger, and R. D. Yates. Age of information aware cache updating with file- and age-dependent update durations. Available on arXiv: 1909.05930.
[18] M. Wang, W. Chen, and A. Ephremides. Reconstruction of counting process in real-time: The freshness of information through queues. In *IEEE ICC*, July 2019.
[19] Y. Sun, Y. Polyanskiy, and E. Uysal-Biyikoglu. Remote estimation of the Wiener process over a channel with random delay. In *IEEE ISIT*, June 2017.
[20] Y. Sun and B. Cyr. Information aging through queues: A mutual information perspective. In *IEEE SPAWC*, June 2018.
[21] J. Chakravorty and A. Mahajan. Remote estimation over a packet-drop channel with Markovian state. Available on arXiv:1807.09706.
[22] A. Arafa and S. Ulukus. Age minimization in energy harvesting communications: Energy-controlled delays. In *Asilomar Conference*, October 2017.
[23] A. Arafa and S. Ulukus. Age-minimal transmission in energy harvesting two-hop networks. In *IEEE Globecom*, December 2017.
[24] X. Wu, J. Yang, and J. Wu. Optimal status update for age of information minimization with an energy harvesting source. *IEEE Transactions on Green Communications and Networking*, 2(1):193–204, March 2018.
[25] A. Arafa, J. Yang, and S. Ulukus. Age-minimal online policies for energy harvesting sensors with random battery recharges. In *IEEE ICC*, May 2018.
[26] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor. Age-minimal online policies for energy harvesting sensors with incremental battery recharges. In *UCSD ITA*, February 2018.
[27] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor. Online timely status updates with erasures for energy harvesting sensors. In *Allerton Conference*, October 2018.
[28] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor. Using erasure feedback for online timely updating with an energy harvesting sensor. In *IEEE ISIT*, July 2019.
[29] S. Farazi, A. G. Klein, and D. R. Brown III. Average age of information for status update systems with an energy harvesting server. In *IEEE Infocom*, April 2018.
[30] S. Leng and A. Yener. Age of information minimization for an energy harvesting cognitive radio. *IEEE Transactions on Cognitive Communications and Networking*, 5(2):427–439, June 2019.
[31] Z. Chen, N. Pappas, E. Bjornson, and E. G. Larsson. Age of information in a multiple access channel with heterogeneous traffic and an energy harvesting node. Available on arXiv: 1903.05066.
[32] R. V. Bhat, R. Vaze, and M. Motani. Throughput maximization with an average age of information constraint in fading channels. Available on arXiv:1911.07499.
[33] J. Ostman, R. Devassy, G. Durisi, and E. Uysal. Peak-age violation guarantees for the transmission of short packets over fading channels. In *IEEE Infocom*, April 2019.
[34] S. Nath, J. Wu, and J. Yang. Optimizing age-of-information and energy efficiency tradeoff for mobile pushing notifications. In *IEEE SPAWC*, July 2017.
[35] Y. Hsu. Age of information: Whittle index for scheduling stochastic arrivals. In *IEEE ISIT*, June 2018.
[36] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano. Scheduling policies for minimizing age of information in broadcast wireless networks. *IEEE/ACM Transactions on Networking*, 26(6):2637–2650, December 2018.
[37] A. Kosta, N. Pappas, A. Ephremides, M. Kountouris, and V. Angelakis. Age and value of information: Non-linear age case. In *IEEE ISIT*, June 2017.
[38] M. Bastopcu and S. Ulukus. Age of information with soft updates. In *Allerton Conference*, October 2018.
[39] M. Bastopcu and S. Ulukus. Minimizing age of information with soft updates. *Journal of Communications and Networks*, 21(3):233–243, June 2019.
[40] B. Buyukates, A. Soysal, and S. Ulukus. Age of information scaling in large networks. In *IEEE ICC*, May 2019.
[41] B. Buyukates, A. Soysal, and S. Ulukus. Age of information scaling in large networks with hierarchical cooperation. In *IEEE Globecom*, December 2019.
[42] J. Zhong, E. Soljanin, and R. D. Yates. Status updates through multicast networks. In *Allerton Conference*, October 2017.
[43] B. Buyukates, A. Soysal, and S. Ulukus. Age of information in two-hop multicast networks. In *Asilomar Conference*, October 2018.
[44] B. Buyukates, A. Soysal, and S. Ulukus. Age of information in multicast networks with multiple update streams. In *Asilomar Conference*, November 2019.
[45] B. Buyukates, A. Soysal, and S. Ulukus. Age of information in multihop multicast networks. *Journal of Communications and Networks*, 21(3):256–267, July 2019.
[46] J. Zhong and R. D. Yates. Timeliness in lossless block coding. In *IEEE DCC*, March 2016.
[47] J. Zhong, R. D. Yates, and E. Soljanin. Timely lossless source coding for randomly arriving symbols. In *IEEE ITW*, November 2018.
[48] P. Mayekar, P. Parag, and H. Tyagi. Optimal lossless source codes for timely updates. In *IEEE ISIT*, June 2018.
[49] D. Ramirez, E. Erkip, and H. V. Poor. Age of information with finite horizon and partial updates. Available on arXiv:1910.00963.
[50] M. Bastopcu, B. Buyukates, and S. Ulukus. Optimal selective encoding for timely updates. In *CISS*, March 2020.
[51] J. Gong, Q. Kuang, X. Chen, and X. Ma. Reducing age-of-information for computation-intensive messages via packet replacement. Available on arXiv: 1901.04654.
[52] B. Buyukates and S. Ulukus. Timely distributed computation with stragglers. Available on arXiv: 1910.03564.
[53] P. Zou, O. Ozel, and S. Subramaniam. Trading off computation with transmission in status update systems. In *IEEE PIMRC*, September 2019.
[54] A. Arafa, R. D. Yates, and H. V. Poor. Timely cloud computing: Preemption and waiting. Available on arXiv:1907.05408.
[55] M. Bastopcu and S. Ulukus. Age of information for updates with distortion. In *IEEE ITW*, August 2019.
[56] M. Bastopcu and S. Ulukus. Age of information for updates with distortion: Constant and age-dependent distortion constraints. Available on arXiv:1912.13493.
[57] M. A. Abd-Elmagid and H. S. Dhillon. Average peak age-of-information minimization in UAV-assisted IoT networks. *IEEE Transactions on Vehicular Technology*, 68(2):2003–2008, February 2019.
[58] J. Liu, X. Wang, and H. Dai. Age-optimal trajectory planning for UAV-assisted data collection. In *IEEE Infocom*, April 2018.
[59] E. T. Ceran, D. Gunduz, and A. Gyorgy. A reinforcement learning approach to age of information in multi-user networks. In *IEEE PIMRC*, September 2018.
[60] H. B. Beytur and E. Uysal-Biyikoglu. Age minimization of multiple flows using reinforcement learning. In *IEEE ICNC*, February 2019.
[61] M. A. Abd-Elmagid, H. S. Dhillon, and N. Pappas. A reinforcement learning framework for optimizing age-of-information in RF-powered communication systems. Available on arXiv: 1908.06367.
[62] E. Cohen and S. Shenker. Replication strategies in unstructured peer-to-peer networks. *ACM SIGCOMM Computer Communication Review*, 32(4):177–190, August 2002.
[63] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.