Age of Information in Relativistic Communication Systems

Antonio Franco, Björn Landfeldt
Department of Electrical and Information Technology
Lund University
Lund, Sweden
e-mail: {antonio.franco, bjorn.landfeldt}@eit.lth.se

***Abstract*—Age of Information (AoI) is a widely used metric to quantify the freshness of updates in communication systems. Existing AoI analyses implicitly assume a shared or synchronized notion of time between transmitter and receiver, thereby neglecting distortions arising from relative motion and gravitational effects. In this paper, we investigate the impact of relativistic time dilation on information freshness and introduce a relativistic formulation of peak Age of Information (pAoI). We first study a special-relativistic setting in which a transmitter moves at constant velocity relative to a receiver, showing that naive timestamp comparison leads to systematic errors in pAoI evaluation, even at moderate velocities. When transmitter velocity is unknown, we characterize the uncertainty induced by velocity estimation and demonstrate that a polynomial estimator can accurately reconstruct the generation time in the receiver reference frame. We then extend the framework to general relativity using a weak-field approximation that jointly accounts for gravitational and kinematic clock-rate distortions. Leveraging established satellite clock modeling techniques, we show that local polynomial approximation enables effective recovery of information freshness despite unknown relativistic effects. Numerical results validate the proposed estimators across all velocity and orbital regimes. This work establishes a first connection between relativistic time modeling and AoI, providing a foundation for freshness analysis in satellite, deep-space, and high-mobility communication systems relevant to space communications and navigation.**

Keywords-age of information; special relativity; general relativity; GNSS; time dilation.

I. Introduction

Timely information delivery is a fundamental requirement in modern communication and control systems, where decisions are often made based on the most recently received status updates. In such systems, the notion of freshness captures how well the information available at a receiver reflects the current state of a remote source. The concept of freshness was formalized through the AoI metric, introduced in [1], which measures the time elapsed since the generation of the latest successfully received update. A closely related measure is the pAoI [2], defined as the maximum age immediately prior to the reception of an update. AoI and pAoI have since become central performance metrics in the analysis and design of communication systems (for a comprehensive survey on AoI, see [3][4]).

Most existing AoI analyses implicitly assume a shared or synchronized notion of time between transmitter and receiver, or equivalently neglect distortions arising from relative motion and gravitational effects. This assumption is reasonable for terrestrial networks, but becomes questionable in high-mobility or space-based systems, where relativistic time dilation can significantly affect the interpretation of timestamps.

In contrast to navigation and timing systems, such as Global Navigation Satellite Systems (GNSS), where relativistic effects are corrected using known orbital parameters [5][6], we consider a communication-theoretic setting in which transmitter motion and gravitational effects are not known a priori and must be inferred from received signals or timestamps. In this context, relativistic distortions directly impact the perceived freshness of information. Motivated by established clock modeling practices in satellite timing systems [7], we adopt a receiver-side polynomial correction framework to mitigate relativistic effects on information freshness. This approach enables accurate reconstruction of the pAoI without requiring explicit knowledge of transmitter velocity or gravitational potential. The main contributions of this paper are as follows:

- We introduce a relativistic formulation of pAoI, explicitly accounting for time dilation induced by relative motion and gravity.
- For special relativity, we show that naive timestamp comparison leads to systematic pAoI errors, and we characterize the uncertainty introduced when transmitter velocity is unknown.
- We propose a carrier-free polynomial estimator that reconstructs the generation time in the receiver reference frame using only timestamped updates, and demonstrate that it consistently outperforms the naive method.
- We extend the estimator framework to general relativity using a weak-field approximation, enabling pAoI recovery under combined gravitational and kinematic clock distortions.
- Through numerical simulations, we validate the proposed estimators and quantify their accuracy improvements over naive AoI evaluation.

We emphasize that the estimators introduced in this work are not designed to reduce the pAoI through scheduling, preemption, or control mechanisms. Instead, they address a measurement problem: how pAoI should be consistently estimated at the receiver when the transmitter and receiver experience different proper times due to relativistic effects. The underlying update generation process and service discipline remain unchanged.

The remainder of this paper is organized as follows. In Section II, related work is presented. In Section III, AoI is characterized in the context of Special Relativity and a pAoI estimator is developed. In Section IV, AoI is characterized in

 ISBN: 978-1-68558-391-0

the context of General Relativity and another more sophisticated pAoI estimator is developed. In Section V, the previous estimators are compared with simulations and results analyzed. Finally, in Section VI, conclusions and future work are discussed.

II. Related Work

A growing body of work has investigated AoI in satellite and non-terrestrial networks, particularly in Low Earth Orbit (LEO) systems [8]–[16]. These contributions analyze AoI and pAoI under stochastic geometry, queueing, and scheduling perspectives, incorporating practical aspects such as intermittent connectivity, propagation delay, interference, and satellite mobility. However, in these works, Doppler effects are typically treated as impairments affecting signal quality or access probability, while propagation delay is often modeled as a deterministic or slot-based quantity. The relativistic nature of time itself – namely, the fact that timestamps generated onboard a moving or gravitationally displaced transmitter correspond to a different proper time than that of the receiver – is not explicitly modeled. Consequently, the AoI metric is evaluated under the implicit assumption that transmitter timestamps are directly comparable to receiver time.

Relativistic effects have been extensively studied in the context of satellite time and frequency synchronization, most notably within GNSSs. The foundational treatment by Ashby [5] provides a comprehensive analysis of both special and general relativistic effects affecting satellite clocks, including gravitational frequency shifts, time dilation due to orbital motion, the Sagnac effect, and relativity of simultaneity. These effects are sufficiently large that, without explicit correction, accurate positioning and timing would be impossible.

Subsequent works have focused on characterizing and modeling residual clock errors after relativistic correction. In [6], the authors analyze periodic variations in Global Positioning System (GPS) satellite clocks arising from orbital dynamics and relativistic effects, demonstrating that residual clock behavior can be effectively modeled using deterministic components. More recently, polynomial clock prediction models have become standard in precise point positioning and timing applications, where satellite clock offsets are represented using low-order polynomials capturing phase, frequency offset, and frequency drift [7].

Receiver-side clock modeling has also been shown to significantly improve synchronization performance. In [17], the authors demonstrate that treating receiver clock offsets as structured stochastic processes rather than white noise yields substantial gains in timing stability, further reinforcing the importance of deterministic clock modeling in satellite timing systems.

To the best of our knowledge, no existing AoI-centric study explicitly incorporates relativistic time dilation or gravitational frequency shifts into the definition, estimation, or statistical characterization of AoI. In particular, the effect of estimating transmitter velocity or gravitational potential on the uncertainty of AoI has not been addressed. This work bridges this gap by integrating relativistic time modeling with AoI analysis, and by proposing estimators that reconstruct the AoI in the receiver's time frame under both special and general relativity.

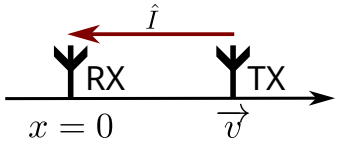

Figure 1. Reference frame for the constant relative speed scenario.

III. Special Relativity

In this section, we develop the framework to model pAoI in a special-relativistic setting. We first introduce the system model for a transmitter moving at constant velocity along the line of sight to the receiver, and show how relativistic time dilation distorts naive timestamp-based pAoI evaluation. We then address the practical case in which the transmitter velocity is unknown at the receiver, characterizing the resulting estimation uncertainty and proposing a timestamp-only polynomial estimator that recovers the generation time in the receiver reference frame.

A. Model

Consider a transmitter (TX) that sends a piece of information $\hat{I}$ to a receiver (RX), where TX always timestamps $\hat{I}$ with its onboard time at the moment of transmission. The RX is located at the origin of its reference frame, while the TX moves at constant speed v along the x-axis (see Figure 1). We assume that a synchronization event has occurred, meaning that TX and RX met at the same position and set their clocks to 0 at that moment. For simplicity, we assume that the physical dimensions of both TX and RX are negligible and that gravitational effects are absent. We further assume that the transmission length of $\hat{I}$ itself – from header to end of content – is negligible, such that the piece of information along with its timestamp, upon reception, is decoded immediately. The propagation medium is vacuum.

We denote the RX proper time – i.e., the time measured in the frame of reference of RX – with t, and the TX proper time with τ. The TX timestamps $\hat{I}$ with τ_g, as measured by its onboard clock. Due to special-relativistic time dilation, the proper time τ of TX and the coordinate time t in the RX frame are related by:

$$\tau = \frac{t}{\gamma}\,, \tag{1}$$

$$\gamma = \frac{1}{\sqrt{1-\beta^2}}\,, \tag{2}$$

$$\beta = \frac{v}{c}\,,$$

where c is the speed of light in vacuum. When TX timestamps τ_g, the RX frame emission time is:

$$t_g = \gamma \tau_g\,. \tag{3}$$

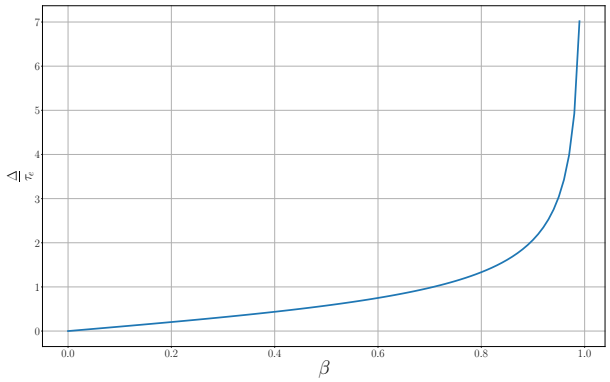

Figure 2. Normalized pAoI at RX vs β.

Since we assume constant speed v, the position of TX at an arbitrary time t is $x(t) = vt$. The update generation position x_g – i.e., the position of TX with respect to RX at the moment of the generation of the update t_g, as measured by the RX onboard clock – is then:

$$x_g = vt_g = v\gamma\tau_g \,. \tag{4}$$

The propagation time is $\frac{x_g}{c}$, and therefore the reception time t_r at RX is:

$$t_r = t_g + \frac{x_g}{c} = \gamma\tau_g + \frac{v\gamma\tau_g}{c} = \gamma(1+\beta)\tau_g \,.$$

Notice that for TX approaching RX (moving toward RX), the sign flips:

$$t_r = \gamma(1-\beta)\tau_g \,.$$

For a more compact representation, we group both multipliers in the Doppler factor:

$$k_\pm = \gamma(1\pm\beta) = \sqrt{\frac{1\pm\beta}{1\mp\beta}} \,, \tag{5}$$

so that

$$t_r = k_\pm \tau_g \,.$$

A naive approach to evaluating the pAoI would be to directly compare the timestamp of $\hat{I}$ with the RX onboard time:

$$\Delta_{\text{naive}} = t_r - \tau_g = \tau_g(k_\pm - 1) \,.$$

However, this yields a negative value for $k_\pm < 1$, which occurs when TX approaches RX. A more correct approach is to calculate the pAoI $\Delta(t_g)$ as:

$$\Delta = t_r - t_g = \frac{|x_g|}{c} = \frac{\gamma|v|\tau_g}{c} = \gamma|\beta|\tau_g \,, \tag{6}$$

where we used (4) to expand x_g. Equation (6) establishes a clear relationship between pAoI and relative speed. In Figure 2, the pAoI, normalized to τ_g, is plotted against β. As shown, the pAoI grows as the speed of TX approaches the speed of light.

Now assume that TX sends updates according to a random process with average rate λ_{TX} updates per second, i.e., with average spacing $\frac{1}{\lambda_{TX}}$. By applying the well-known formula for the relativistic Doppler effect (similar to (1)) the arrival rate as seen by RX is:

$$\lambda_{RX} = \frac{\lambda_{TX}}{\gamma} \,,$$

More generally, let θ denote the angle between the TX velocity and the Line of Sight (LOS) toward RX. Then:

$$\lambda_{RX} = \frac{\lambda_{TX}}{\gamma(1-\beta\cos(\theta))} \,.$$

This shows that updates arrive faster at the RX when the TX moves toward it (resulting in lower AoI measured onboard the RX), and slower when it moves away (resulting in higher AoI measured onboard the RX). Note that the entire AoI process is affected by the Doppler effect, which can yield long tails in the distribution even when this effect is unintended.

B. Estimating pAoI

Until now, we have assumed that the magnitude and direction of the transmitter velocity are known, which allows us to recover the AoI in RX time via (6). We now study how uncertainty in v affects the AoI measured on the RX side.

1) Carrier Frequency Estimation: The simplest method to estimate v is for TX and RX to share knowledge of a carrier frequency f_0. Then RX can simply calculate the difference with the measured frequency $\hat{f}_0$ and from there determine the Doppler factor, and thus t_g in its reference frame via (1). In particular, we assume that TX moves inertially at constant speed along the LOS, that no scattering occurs, and that the physical layer decodes correctly. We ignore symbol deformation and similar effects. In this case, the pAoI at reception time is itself an estimate:

$$\Delta = t_r - \hat{t}_g \,, \tag{7}$$

where $\hat{t}_g$ is RX's estimate of the generation time expressed in the RX time coordinate. Since γ is a function of β (2), we have:

$$\hat{t}_g = \gamma(\hat{\beta})\tau_g \,.$$

We define the conversion-error random variable as:

$$\epsilon = t_g - \hat{t}_g = \gamma(\beta)\tau_g - \gamma(\hat{\beta})\tau_g \,.$$

The mapping $\gamma(\hat{\beta})$ can be approximated by a first-order Taylor expansion around the true β. The first-order approximation is justified under standard asymptotic conditions for Doppler-based estimators, such as sufficiently high Signal-to-Noise Ratio (SNR) or long coherent observation time [18]. Thus:

$$\begin{aligned}\gamma(\hat{\beta}) &= \gamma(\beta) + \gamma'(\beta)\,(\hat{\beta}-\beta) + O\Big((\hat{\beta}-\beta)^2\Big) \\ &\approx \gamma(\beta) + \gamma'(\beta)\,(\hat{\beta}-\beta) \,,\end{aligned}$$

which yields:

$$\epsilon \approx -\tau_g\gamma'(\beta)\,(\hat{\beta}-\beta) \,.$$

Finally, since in (7) the only uncertainty is in the estimated t_g, we can conclude:

$$\mathrm{Var}(\Delta) \approx \mathrm{Var}(\epsilon) = \mathrm{Var}\left(-\tau_g\gamma'(\beta)\,(\hat{\beta}-\beta)\right)$$

$$= \tau_g^2 \gamma'(\beta)^2 \operatorname{Var}(\hat{\beta} - \beta) = \tau_g^2 \beta^2 \gamma(\beta)^3 \operatorname{Var}(\hat{\beta}) \,, \qquad (8)$$

where we used (2) to expand $\gamma'(\beta)$.

It remains to relate $\operatorname{Var}(\hat{\beta})$ to $\operatorname{Var}(\hat{f}_r) = \sigma_f^2$, where $\hat{f}_r$ is the estimate of the Doppler-shifted frequency as seen from RX. By solving (5) for β in terms of $k_\pm$, we obtain:

$$\beta_- = \frac{1 - k_-^2}{1 + k_-^2}$$
$$\beta_+ = -\frac{1 - k_+^2}{1 + k_+^2} \,.$$

We proceed with the case k_-; the derivation for k_+ is analogous and yields the same result. Taking the first derivative with respect to β we get:

$$\gamma'(\beta) = \beta \gamma^3 \,.$$

Using the same logic applied in deriving (8), and noting that the Doppler factor is also the ratio between f_r and f_0, we obtain:

$$\operatorname{Var}(\hat{\beta}) = \operatorname{Var}(\beta(\hat{k}_-)) \approx \operatorname{Var}(\beta(k_-) + \beta'(k_-)(\hat{k}_- - k_-))$$
$$= \beta'(k_-)^2 \operatorname{Var}(\hat{k}_-)$$

Finally, substituting this result into (8) yields:

$$\operatorname{Var}(\Delta) \approx \tau_g^2 \beta^2 \gamma^6 \frac{16 k_-^2}{(1 + k_-^2)^4} \frac{\sigma_f^2}{f_0^2} \,.$$

Thus, the standard deviation in the estimated pAoI grows proportionally with the uncertainty in the carrier frequency estimate.

2) pAoI Polynomial Estimator: The receiver observes a stream of timestamped updates. Each update carries the transmitter-side generation timestamp $\tau_g^{(n)}$, expressed in the TX proper-time reference frame, and is received at the RX at time $t_r^{(n)}$, measured by the RX onboard clock. While $\tau_g^{(n)}$ is explicitly available in the packet payload, the corresponding generation time expressed in the RX time coordinate, $t_g^{(n)}$, is not directly observable due to relativistic time dilation and propagation effects. As a result, a naive AoI computation based on directly comparing $t_r^{(n)}$ and $\tau_g^{(n)}$ yields incorrect or even unphysical results, as discussed in Section III-A.

The objective of the pAoI polynomial estimator is therefore to reconstruct, at the receiver, an estimate of the generation time expressed in the RX time coordinate, $\hat{t}_g(\tau_g^{(n)})$, using only the sequence of observed timestamp/arrival-time pairs $\{(\tau_g^{(i)}, t_r^{(i)})\}$. Once such an estimate is available, the pAoI at reception can be formed as

$$\hat{\Delta}_n \triangleq t_r^{(n)} - \hat{t}_g(\tau_g^{(n)}),$$

without relying on carrier-based Doppler estimation or prior knowledge of the transmitter velocity.

Polynomial approximation of the time-transfer mapping: Motivated by established clock modeling techniques in high-precision timing systems [7], we approximate the mapping between the TX timestamp τ and the RX reception time t_r locally in time by a low-order polynomial. Over a sufficiently short observation window, we write

$$\hat{t}_r(\tau) = a_0 + a_1 \tau + a_2 \tau^2 + \dots \,, \qquad (9)$$

where the coefficients $\{a_k\}$ capture the combined effects of relativistic time dilation, propagation delay, and any slowly varying mismatch between the TX and RX clocks.

Least-squares coefficient estimation: Given a sliding window of N received updates $\{(\tau_g^{(i)}, t_r^{(i)})\}_{i=n-N+1}^{n}$, the RX estimates the polynomial coefficients by Ordinary Least Squares (OLS). Defining the observation vector $\mathbf{t}_r$ and the Vandermonde design matrix $\mathbf{X}$ as:

$$\mathbf{t}_r \triangleq \begin{bmatrix} t_r^{(n-N+1)} & \dots & t_r^{(n)} \end{bmatrix}^T ,$$
$$\mathbf{X} \triangleq \begin{bmatrix} 1 & \tau_g^{(n-N+1)} & (\tau_g^{(n-N+1)})^2 & \dots \\ \vdots & \vdots & \vdots & \\ 1 & \tau_g^{(n)} & (\tau_g^{(n)})^2 & \dots \end{bmatrix} ,$$

the coefficient estimate is:

$$\hat{\mathbf{a}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{t}_r \,.$$

The matrix $\mathbf{X}^T\mathbf{X}$ is invertible provided that the RX observes at least $p{+}1$ updates at distinct timestamps over the estimation window, which is naturally satisfied in any non-degenerate AoI process.

From polynomial fit to generation-time estimation: In the constant-velocity, LOS scenario considered in this section, special relativity implies that the reception time is related to the TX timestamp by a Doppler factor $k_\pm$, i.e., $t_r = k_\pm \tau$. Consequently, the Doppler factor can be identified with the derivative of the reception-time mapping with respect to τ. Using the fitted polynomial (9), the RX forms

$$\hat{k}_n \triangleq \frac{d}{d\tau} \hat{t}_r(\tau)\big|_{\tau = \tau_g^{(n)}} .$$

The corresponding velocity and Lorentz factor estimates follow directly from the relativistic Doppler relation, yielding $\hat{\beta}_n$ and $\hat{\gamma}_n$. Finally, the RX estimates the generation time expressed in its own reference frame as

$$\hat{t}_g(\tau_g^{(n)}) = \hat{\gamma}_n \, \tau_g^{(n)} .$$

pAoI estimation procedure: The complete AoI synchronization and estimation procedure is summarized in Figure 3. The algorithm operates purely on timestamped packet arrivals and produces a pAoI estimate at each reception instant once the estimator reaches steady state.

IV. General Relativity

In contrast to the special-relativistic setting, where the relationship between transmitted timestamps and reception times admits a closed-form expression, the general-relativistic case does not generally yield an explicit mapping, like the one expressed by (3). Relativistic clock-rate differences and propagation effects depend on the gravitational field and on the spacetime geometry encountered along the signal path. As

Require: Polynomial degree p; window size $N \geq p+1$; stream of received updates $\{(\tau_g^{(n)}, t_r^{(n)})\}$

Ensure: pAoI estimate at reception $\hat{\Delta}_n$ and generation-time estimate $\hat{t}_g(\tau_g^{(n)})$

1: Maintain a sliding window $\mathcal{W}_n = \{(\tau_g^{(i)}, t_r^{(i)})\}_{i=n-N+1}^{n}$

2: **if** $|\mathcal{W}_n| < N$ **then**

3: **return** (insufficient samples)

4: Build $\mathbf{X} \in \mathbb{R}^{N\times(p+1)}$ with rows $\left[1,\ \tau_g^{(i)},\ (\tau_g^{(i)})^2,\ \ldots,\ (\tau_g^{(i)})^p\right]$

5: Stack $\mathbf{t}_r \leftarrow [t_r^{(n-N+1)}, \ldots, t_r^{(n)}]^T$

6: Compute OLS coefficients $\hat{\mathbf{a}} \leftarrow (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{t}_r$

7: Evaluate Doppler-factor estimate via the derivative:

$$\hat{k}_n \leftarrow \frac{d}{d\tau}\hat{t}_r(\tau)\big|_{\tau=\tau_g^{(n)}} = \sum_{m=1}^{p} m\,\hat{a}_m \left(\tau_g^{(n)}\right)^{m-1}$$

8: Invert the SR Doppler relation (LOS case) to obtain

$$\hat{\beta}_n \leftarrow \frac{\hat{k}_n^2 - 1}{\hat{k}_n^2 + 1}, \qquad \hat{\gamma}_n \leftarrow \frac{1}{\sqrt{1-\hat{\beta}_n^2}}$$

9: Compute generation-time estimate in RX coordinates:

$$\hat{t}_g(\tau_g^{(n)}) \leftarrow \hat{\gamma}_n\,\tau_g^{(n)}$$

10: Output pAoI estimate at reception:

$$\hat{\Delta}_n \leftarrow t_r^{(n)} - \hat{t}_g(\tau_g^{(n)})$$

Figure 3. pAoI Polynomial Synchronization at RX (timestamp-only).

a result, the AoI evaluation problem in the general-relativistic setting naturally leads to an estimator-based formulation, which reconstructs the time-transfer mapping directly from timestamped observations.

A. Model

Consider a TX that sends timestamped updates to an RX, where each update is labeled with the TX onboard clock reading at the moment of generation. Both TX and RX are equipped with ideal clocks measuring their respective proper times. Unlike the special-relativistic case, we now allow the spacetime to be curved due to the presence of a gravitational field. TX and RX follow prescribed worldlines in a weak gravitational field, and their clocks generally tick at different rates due to both relative motion and gravitational time dilation. We assume that a synchronization event has occurred at some initial spacetime point, after which TX and RX clocks are related only through the exchange of timestamped signals. The physical dimensions of TX and RX are assumed negligible, and the transmission duration of each update is assumed small compared to the timescales of interest. Signal propagation occurs through vacuum along null trajectories of the spacetime geometry (a null trajectory is the spacetime path followed by light or radio signals; while such signals always propagate at the speed of light locally, gravity can bend their paths and modify their travel times, an effect accounted for by general relativity).

Let τ denote the TX proper time and t the RX proper time, measured by their respective onboard clocks. Each update generated at TX at time τ_g is received by RX at time t_r. In the presence of gravitational fields, the relationship between τ_g and t_r is governed by relativistic clock-rate differences and by signal propagation along curved spacetime trajectories. These effects jointly induce a deterministic but generally unknown time-transfer mapping between TX timestamps and RX reception times, which we denote by:

$$t_r = \mathcal{T}(\tau_g). \tag{10}$$

Unlike the special-relativistic case, this mapping does not generally admit a closed-form expression and may vary slowly over time due to changes in gravitational potential or non-inertial motion. For reference, the usual weak-field approximation for the rate at which TX proper time τ elapses with respect to RX time t [5] is $\frac{d\tau}{dt} \approx 1 - \frac{\Phi(x)}{c^2} - \frac{v^2}{2c^2}$, where $\Phi(x)$ is the Newtonian gravitational potential at position x and v is the instant relative velocity.

The pAoI at reception is defined with respect to RX time as

$$\Delta = t_r - t_g\,,$$

where t_g denotes the generation time expressed in the RX time coordinate. Since t_g is not directly observable at the receiver, evaluating the AoI in the general-relativistic setting requires estimating the inverse relationship between the transmitted timestamp τ_g and the corresponding generation time in the RX reference frame. This observation motivates the estimator-based approach developed in the following subsection.

B. pAoI Polynomial Estimator

The receiver observes a sequence of timestamped updates, where each update carries the TX proper-time timestamp $\tau_g^{(n)}$ and is received at RX time $t_r^{(n)}$. As established in Section IV-A (10), these quantities are related through an unknown time-transfer mapping $t_r = \mathcal{T}(\tau_g)$, which captures both relativistic clock-rate differences and signal propagation effects. The objective of the estimator is to reconstruct, at the receiver, an estimate of the generation time expressed in the RX time coordinate, $\hat{t}_g(\tau_g^{(n)})$, using only the observed timestamp/arrival-time pairs. Once such an estimate is available, the pAoI at reception is formed as

$$\hat{\Delta}_n \triangleq t_r^{(n)} - \hat{t}_g(\tau_g^{(n)}).$$

Local polynomial approximation of the time-transfer mapping: Under the weak-field and slow-variation assumptions, the time-transfer mapping $\mathcal{T}(\cdot)$ is smooth and can be locally approximated over a finite observation window by a low-order polynomial [7]. Specifically, we model

$$\hat{t}_r(\tau) = a_0 + a_1\tau + a_2\tau^2 + \ldots + a_p\tau^p. \tag{11}$$

Coefficient estimation via least squares: Given a sliding window of N received updates $\{(\tau_g^{(i)}, t_r^{(i)})\}_{i=n-N+1}^{n}$, the

Require: Polynomial degree p; window size $N \geq p+1$; stream of received updates $\{(\tau_g^{(n)}, t_r^{(n)})\}$
Ensure: pAoI estimate $\hat{\Delta}_n$ at each reception
1: Initialize $\hat{t}_g^{(0)} \leftarrow 0$
2: Maintain a sliding window $\mathcal{W}_n = \{(\tau_g^{(i)}, t_r^{(i)})\}_{i=n-N+1}^{n}$
3: **if** $|\mathcal{W}_n| < N$ **then**
4: **return** (insufficient samples)
5: Construct Vandermonde matrix $\mathbf{X}$ from timestamps $\tau_g^{(i)}$
6: Stack observation vector $\mathbf{t}_r = [t_r^{(n-N+1)}, \ldots, t_r^{(n)}]^T$
7: Compute polynomial coefficients $\hat{\mathbf{a}} \leftarrow (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{t}_r$
8: Evaluate local time-transfer rate:

$$\hat{\rho}_n \leftarrow \sum_{m=1}^{p} m\,\hat{a}_m \left(\tau_g^{(n)}\right)^{m-1}$$

9: Update generation-time estimate:

$$\hat{t}_g^{(n)} \leftarrow \hat{t}_g^{(n-1)} + \hat{\rho}_n\left(\tau_g^{(n)} - \tau_g^{(n-1)}\right)$$

10: Compute pAoI estimate:

$$\hat{\Delta}_n \leftarrow t_r^{(n)} - \hat{t}_g^{(n)}$$

Figure 4. pAoI Polynomial Estimation under General Relativity.

receiver estimates the polynomial coefficients by OLS. Defining the observation vector $\mathbf{t}_r$ and the Vandermonde design matrix $\mathbf{X}$ constructed from the timestamps $\tau_g^{(i)}$, the coefficient estimate is:

$$\hat{\mathbf{a}} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{t}_r,$$

provided that $N \geq p+1$ and that the timestamps in the window are distinct.

Reconstruction of the generation-time mapping: The derivative of the fitted polynomial yields an estimate of the local rate at which RX time evolves with respect to TX proper time:

$$\hat{\rho}_n \triangleq \frac{d}{d\tau}\hat{t}_r(\tau)\bigg|_{\tau=\tau_g^{(n)}} = \sum_{m=1}^{p} m\,\hat{a}_m \left(\tau_g^{(n)}\right)^{m-1}.$$

This quantity represents the instantaneous time-transfer rate between TX timestamps and RX time, incorporating all relativistic and propagation effects. The receiver reconstructs the generation time expressed in the RX time coordinate by integrating the estimated local rate. In discrete time, this is performed recursively as

$$\hat{t}_g^{(n)} = \hat{t}_g^{(n-1)} + \hat{\rho}_n\left(\tau_g^{(n)} - \tau_g^{(n-1)}\right),$$

with initialization $\hat{t}_g^{(0)} = 0$ at the synchronization event.

pAoI estimation: Finally, the pAoI at reception time $t_r^{(n)}$ is estimated as

$$\hat{\Delta}_n = t_r^{(n)} - \hat{t}_g^{(n)}.$$

The complete estimation procedure is summarized in Figure 4.

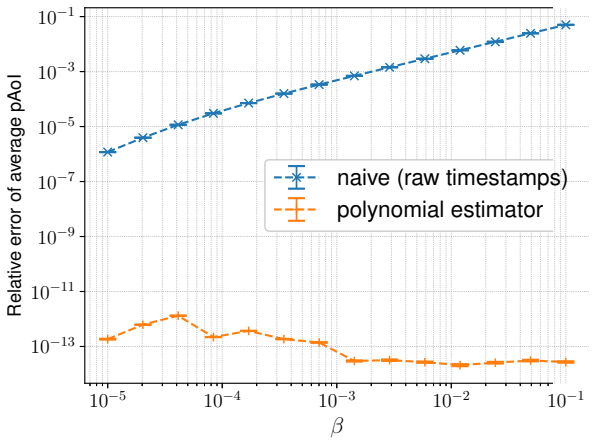

Figure 5. Estimated pAoI by Figure 3 vs Naive relative error.

TABLE I. SIMULATION PARAMETERS FOR THE SPECIAL-RELATIVISTIC CONSTANT-VELOCITY SCENARIO.

Parameter	Value
Polynomial degree p	2
Update generation rate λ_{TX}	2 Hz
Observation window size N	120 samples
Update generation process	Poisson process
Timestamp noise std. σ_t	1 ns

V. NUMERICAL RESULTS

In this section, we present our simulation studies. All plots involving simulations include a sufficient warm-up period before measurements are taken. Confidence intervals are shown, except where they are too tight to show at 95% confidence.

To evaluate estimator performance across relativistic regimes, we simulate the special-relativistic constant-velocity scenario introduced in Section III. The transmitter generates timestamped updates according to a Poisson process in its proper time with average rate λ_{TX}, and each update carries its generation timestamp τ_g. We chose to use a Poissonian generation process as it is memoryless, thus modeling independent single transmissions; the estimator itself is generation process agnostic, as long as the process is stationary and ergodic. Updates propagate over a LOS channel, and the receiver observes the corresponding reception times t_r in its own reference frame, distorted by relativistic time dilation and corrupted by Additive White Gaussian timestamp Noise (AWGN) with standard deviation σ_{t_r}.

We compare two receiver-side pAoI estimation strategies: (i) a naive approach that directly interprets transmitted timestamps in the receiver frame, and (ii) the polynomial AoI estimator described in Figure 3. To the best of our knowledge, no existing work addresses the estimation of AoI under relativistic time distortion. The polynomial estimator operates on a sliding window of N received updates and locally approximates the mapping between transmitted timestamps and reception times

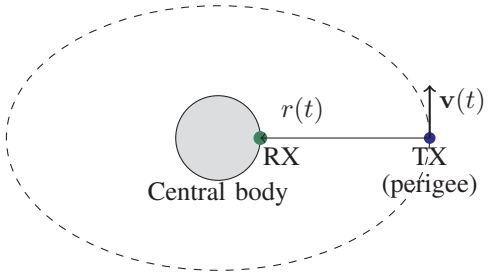

Figure 6. Simulation geometry for the general-relativistic orbital scenario.

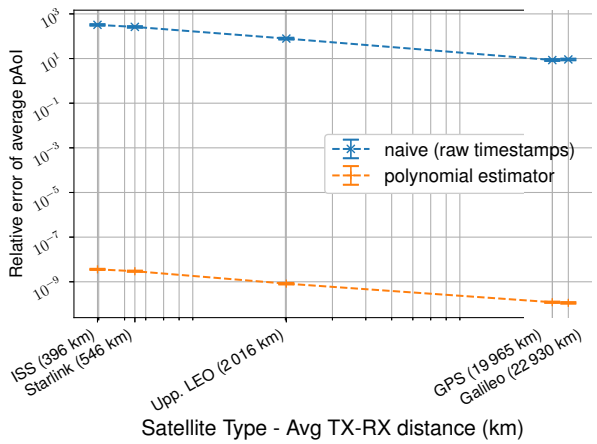

Figure 7. Estimated pAoI by the Algorithm in Figure 4 vs Naive relative error under general-relativistic time distortion. pAoI is presented vs the average altitude of TX. The most common type of satellite at that altitude is presented as well as a label.

by a polynomial of degree p. From this approximation, the receiver reconstructs the generation time expressed in its own reference frame and computes the estimated pAoI by comparing it with the reception time.

The normalized transmitter speed β is swept logarithmically from values representative of GNSS systems to increasingly relativistic regimes. For each β, the simulation is allowed to reach steady state before collecting measurements. The parameters used in this numerical study are reported in Table I.

Figure 5 compares the true average pAoI with estimates obtained using the naive timestamp comparison and the proposed polynomial estimator. The metric used is the relative error of the average pAoI with respect to the true average pAoI, i.e.,

$$\frac{|\mathbb{E}[\hat{\Delta}] - \mathbb{E}[\Delta]|}{\mathbb{E}[\Delta]}.$$

Notice that said error is dimensionless.

The results show that the naive approach consistently overestimates or underestimates pAoI depending on the relative motion, while the estimator closely tracks the true value. Naive AoI evaluation fails to account for time dilation, whereas the proposed estimator successfully recovers the correct pAoI by exploiting temporal structure across updates.

To test the pAoI estimator described in Figure 4, we consider a concrete instantiation of the general-relativistic model introduced in Section IV-A, corresponding to a transmitter moving on a Keplerian geosynchronous orbit around Earth as the central body (Figure 6). Note that RX and TX are assumed to be always in direct LOS, being TX in geosynchronous orbit around Earth, thus the distance separating them is the instantaneous radius of TX orbit minus the radius of the Earth, i.e., the minimum euclidean distance between them. No atmospheric effects are introduced – i.e., the transmission medium is vacuum; signal propagation occurs through vacuum along null trajectories of the spacetime geometry.

TABLE II. SIMULATION PARAMETERS FOR THE GENERAL-RELATIVISTIC ORBITAL SCENARIO.

Parameter	**Value**
Gravitational parameter μ	3.986×10^{14} m^3/s^2
Central radius R	6 371 km
Orbital eccentricity e	0.01
Perigee speed ratio $\beta_p = v_p/c$	see Table III
Update generation rate λ_{TX}	2 Hz
Update generation process	Poisson process
Observation window size N	200 samples
Polynomial degree	3
Timestamp noise std. σ_t	1 ns

TABLE III. ORBITAL REGIMES MAPPING β_p TO PERIGEE (h_p) AND APOGEE (h_a) ALTITUDE FOR THE SIMULATION IN FIGURE 7.

Typical usage	β_p	h_p [km]	h_a [km]
ISS	2.585×10^{-5}	332.0	467.4
Starlink	2.557×10^{-5}	480.4	618.9
Upper LEO	2.325×10^{-5}	1 915.9	2 083.3
GPS GNSS	1.305×10^{-5}	19 933.0	20 464.4
Galileo GNSS	1.237×10^{-5}	22 924.6	23 516.4

A simulation run is initialized with the following parameters:

- μ: gravitational parameter of Earth, i.e., the product of the gravitational constant G and the mass of Earth (RX and TX masses are considered negligible with respect to the mass of the Earth).
- R: radius of the Earth: $R \approx 6\,371$ km.
- $e \in [0, 1)$: orbital eccentricity of the transmitter, i.e., how much the orbit deviates from a perfect circle.
- β_p: normalized perigee speed, i.e., $\beta_p = v_p/c$, i.e., the fastest orbital speed TX achieves when at the perigee; along with the previous parameters it is sufficient to completely define TX orbit at any point in relative time. Note that, given μ, e and R, the normalized perigee speed is upper bounded, i.e., $0 \leq \beta_p \leq \frac{1}{c}\sqrt{\frac{\mu(1+e)}{R}} \approx 2.8 \times 10^{-5}$.
- λ_{TX}: update generation rate in the transmitter proper time. As per the Special Relativity case, we chose to generate updates according to a Poissonian Process.
- σ_{t_r}: standard deviation of additive reception timestamp noise at the receiver.
- p: polynomial degree used by the estimator.
- N: estimator window size (number of most recent updates

used for fitting).

The input parameter β_p completely defines the instantaneous altitude of TX, the complete mapping is presented in Table III, ranging from the International Space Station (ISS) to Galileo GNSS altitude.

The life cycle of a generated update goes through the following steps:

1) The update is generated at time τ_n – as measured by the clock onboard TX – and timestamped.
2) Kepler's equation is solved to find the instantaneous eccentric anomaly $E(\tau_n)$:

$$E(\tau_n) - e\sin\left(E(\tau_n)\right) = \tau_n\sqrt{\frac{\mu}{a^3}}\,,$$

where a is the semi-major axis of TX orbit:

$$a = \frac{\mu(1+e)}{v_p^2(1-e)}\,, \tag{12}$$

and v_p is the perigee speed:

$$v_p = \beta_p c\,.$$

3) The instantaneous orbital distance $r(\tau_n)$ is then obtained as:

$$r(\tau_n) = a\big(1 - e\cos\left(E(\tau_n)\right)\big) - R\,.$$

4) From the previous, the instantaneous orbital speed is computed using the vis-viva equation:

$$v(\tau_n) = \sqrt{\mu\left(\frac{2}{r(\tau_n)} - \frac{1}{a}\right)}\,.$$

5) The receiver is assumed to be inertial and located at a fixed position on the surface of the Earth, defining the reference frame in which the pAoI is evaluated. Gravitational effects are modeled through a static Newtonian potential:

$$\Phi(r) = -\mu/r\,, \tag{13}$$

where μ is the standard gravitational parameter of Earth (the central body) and r denotes the instantaneous TX-RX distance. Under the weak-field, slow-motion approximation, and the fact that we assumed continuous LOS, the mapping between proper time and receiver coordinate time is obtained by integrating the rate [5]:

$$\frac{dt}{d\tau} \approx 1 - \frac{\Phi(r)}{c^2} - \frac{v^2}{2c^2}\,.$$

By substituting (13) in the previous we finally obtain the time dilation rate:

$$\frac{dt}{d\tau} \approx 1 + \frac{\mu}{r(\tau_n)c^2} - \frac{v(\tau_n)^2}{2c^2}\,, \tag{14}$$

that, integrated, gives us the true mapping function $\mathcal{T}(\tau_g)$ (10), used to calculate the simulated true reception time $t_g^{(n)}$.

6) The propagation delay is modeled as a LOS radial delay, as we always assume LOS. We also add AWGN $w_n \sim \mathcal{N}(0, \sigma_{t_r}^2)$; the time observed at reception in the RX frame of reference becomes:

$$\hat{t}_r^{(n)} = t_g^{(n)} + \frac{r(\tau_n)}{c} + w_n\,.$$

7) Finally RX, adds $\hat{t}_r^{(n)}$ to the observation window, and uses Figure 4 – when there are sufficient observations, i.e., $n \geq N \geq p+1$ – to produce the current estimation of the true pAoI $\hat{\Delta}(\tau_n)$.

To ensure numerical stability of the least-squares fitting procedure, the transmitter timestamps τ_g within each window are centered and normalized prior to constructing the polynomial regression matrix. This avoids the severe ill-conditioning associated with Vandermonde matrices built on unscaled time variables [19] and ensures robust coefficient estimation even for higher-order polynomials. To avoid initialization bias, the estimator is explicitly anchored when the first complete observation window becomes available by initializing the estimated generation time to the corresponding propagation-corrected observation. Subsequent estimates are obtained recursively by integrating the estimated local clock rate. This approach prevents spurious offsets due to estimator warm-up and ensures that the estimated pAoI does not suffer from numerical artifacts.

In Figure 7, we plotted the estimated pAoI by Figure 4 vs true AoI for different type of satellites (see Table III). The parameters used in the numerical study are described in Table II. The estimator outperforms the naive approach by several orders of magnitude in all regimes. Note that, due to (12), there is a square inverse proportionality between the average altitude of TX and the perigee speed.

The absolute error levels in Figures 5 and 7 are not directly comparable: the special-relativistic mapping is linear and closed-form, yielding a well-conditioned problem, whereas the general-relativistic mapping is nonlinear and time-varying, requiring local approximation and numerical integration under significantly more challenging conditions. Additionally, note that the naive estimator exhibits markedly different behavior in the two scenarios. Specifically, constant distortions under special relativity induce only a linear bias, leading to bounded estimation errors; in contrast, time-varying distortions under general relativity are integrated over the observation horizon, leading to cumulative errors.

VI. Conclusion and Future Work

This paper established a relativistic perspective on AoI, showing that information freshness is fundamentally affected by time dilation due to relative motion and gravity. We demonstrated that naive timestamp comparison leads to systematic errors in pAoI evaluation, even at low relative speeds. To address this, we proposed polynomial estimators that reconstruct the generation time in the receiver reference frame using only timestamped updates. For special relativity, the proposed estimator consistently outperforms the naive method across all velocity regimes. The framework was further extended to general relativity using a weak-field approximation, enabling accurate pAoI recovery under combined gravitational and kinematic

distortions. Numerical results validate the effectiveness of the proposed approach and highlight the importance of relativistic modeling in high-mobility and space-based communication systems.

Several directions emerge naturally from this work. First, while the present analysis relies on weak-field approximations, extending the framework to strong gravitational fields would enable the study of information freshness in more extreme environments, such as deep-space missions or communication near compact astrophysical objects. Second, the general-relativistic model considered here assumes non-relativistic orbital velocities; incorporating regimes in which velocities approach the speed of light would require abandoning the weak-field expansion and adopting a fully relativistic spacetime description, potentially revealing qualitatively new AoI behaviors. Finally, this paper focused primarily on mean AoI and pAoI. An important extension is the characterization of the full AoI distribution under relativistic time distortions, which would allow the analysis of threshold-based freshness violations and quality-of-service guarantees. Such a distributional perspective is particularly relevant for safety-critical and real-time systems, where rare but severe AoI excursions can be as important as average performance.

ACKNOWLEDGMENT

This work was sponsored in part by the Excellence Center at Linköping-Lund in Information Technology (ELLIIT) Sweden.

REFERENCES

[1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?", in *Proc. 2012 IEEE INFOCOM*, Mar. 2012, pp. 2731–2735. DOI: 10.1109/INFCOM.2012.6195689

[2] M. Costa, M. Codreanu, and A. Ephremides, "Age of information with packet management", in *Information Theory (ISIT), 2014 IEEE International Symposium on*, Jun. 2014, pp. 1583–1587. DOI: 10.1109/ISIT.2014.6875100

[3] A. Kosta, N. Pappas, and V. Angelakis, "Age of information: A new concept, metric, and tool", *Foundations and Trends in Networking*, vol. 12, no. 3, pp. 162–259, 2017.

[4] N. Pappas, M. A. Abd-Elmagid, B. Zhou, W. Saad, and H. S. Dhillon, *Age of Information: Foundations and Applications*. Cambridge University Press, 2023, ISBN: 9781108950275.

[5] N. Ashby, "Relativity in the global positioning system", *Living Reviews in Relativity*, vol. 6, no. 1, pp. 1–45, Jan. 2003, ISSN: 1433-8351. DOI: 10.12942/lrr-2003-1

[6] K. L. Senior, J. R. Ray, and R. L. Beard, "Characterization of periodic variations in the gps satellite clocks", *GPS solutions*, vol. 12, no. 3, pp. 211–225, 2008.

[7] H. Ge, B. Li, T. Wu, and S. Jiang, "Prediction models of gnss satellite clock errors: Evaluation and application in ppp", *Advances in Space Research*, vol. 68, no. 6, pp. 2470–2487, 2021.

[8] Z. Gao, A. Liu, C. Han, and X. Liang, "Non-orthogonal multiple access-based average age of information minimization in leo satellite-terrestrial integrated networks", *IEEE Transactions on Green Communications and Networking*, vol. 6, no. 3, pp. 1793–1805, 2022. DOI: 10.1109/TGCN.2022.3159559

[9] Y. Li, Y. Xu, Q. Zhang, and Z. Yang, "Age-optimized multihop information update mechanism on the leo satellite constellation via continuous time-varying graphs", *IEEE Internet of Things Journal*, vol. 10, no. 8, pp. 7189–7203, 2023. DOI: 10.1109/JIOT.2022.3229028

[10] C.-Y. Lin and W. Liao, "Energy-aware age of information (aoi) minimization for internet of things in noma-based leo satellite networks", in *2024 IEEE 99th Vehicular Technology Conference (VTC2024-Spring)*, 2024, pp. 1–5. DOI: 10.1109/VTC2024-Spring62846.2024.10683037

[11] D. Zhang et al., "Age of information variation of leo satellite-terrestrial uplink transmissions", *IEEE Transactions on Vehicular Technology*, vol. 74, no. 6, pp. 9645–9655, 2025. DOI: 10.1109/TVT.2025.3540572

[12] H. Zarini, S. M. Kazemi, M. Sookhak, E. Uysal, and S. Chatzinotas, "Age of information in leo satellite communications supported by bd-ris", in *ICC 2025 - IEEE International Conference on Communications*, 2025, pp. 1–6. DOI: 10.1109/ICC52391.2025.11161975

[13] B. Gabr and M. A. Kishk, "Statistical analysis for average peak age of information in leo satellite-enabled iot", in *2025 IEEE 101st Vehicular Technology Conference (VTC2025-Spring)*, 2025, pp. 1–6. DOI: 10.1109/VTC2025-Spring65109.2025.11174394

[14] M. Liao et al., "Multihop age-of-information-enhanced routing mechanism for delay-sensitive service in the leo satellite networks", *IEEE Internet of Things Journal*, vol. 12, no. 23, pp. 51 168–51 181, 2025. DOI: 10.1109/JIOT.2025.3612043

[15] B. Gabr and M. A. Kishk, "Performance analysis of average peak age of information in leo satellite-enabled iot networks", *IEEE Transactions on Wireless Communications*, vol. 25, pp. 5186–5199, 2026. DOI: 10.1109/TWC.2025.3616774

[16] L. Badia and A. Munari, "Satellite intermittent connectivity and its impact on age of information for finite horizon scheduling", in *2025 12th Advanced Satellite Multimedia Systems Conference and the 18th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, 2025, pp. 1–8. DOI: 10.1109/ASMS/SPSC64465.2025.10946056

[17] W. Qin et al., "The benefits of receiver clock modelling in satellite timing", *Sensors*, vol. 21, no. 2, 2021, ISSN: 1424-8220. DOI: 10.3390/s21020466

[18] D. C. Rife and R. R. Boorstyn, "Single tone parameter estimation from discrete-time observations", *IEEE Transactions on Information Theory*, vol. 20, no. 5, pp. 591–598, 1974.

[19] V. Y. Pan, "How bad are vandermonde matrices?", *SIAM Journal on Matrix Analysis and Applications*, vol. 37, no. 2, pp. 676–694, 2016.