To Share or Not to Share: A Threshold Analysis of AoI under Server Compromise

Alessandro Buratto*, Andrea Munari†, Estefania Recayte†, and Leonardo Badia*
*Dept. Information Engineering, University of Padova, 35131 Padua, Italy
†Institute of Communications and Navigation, German Aerospace Center (DLR), Weßling, Germany
Email: alessandro.buratto.1@studenti.unipd.it, estefania.recayte@dlr.de, andrea.munari@dlr.de leonardo.badia@unipd.it

***Abstract*—We consider a scenario where a source is sending sensitive information, under a general aim of information freshness, e.g., for real-time applications. This objective is quantified through age of information (AoI), for which we focus on the core formulas of queueing theory with preemption. While representing an idealization, this approach is chosen for its analytical tractability, which allows for the derivation of closed-form expressions that make the underlying trade-offs transparent. However, we consider that the source can use two servers, even in parallel, one of which is slower but trusted, whereas the other is faster but potentially compromised and insecure. Our analysis determines how much of potentially sensitive information must be sent to the insecure server under a multi-objective combination of keeping the average AoI low, while at the same time avoiding to disclose too much information. Our analytical findings, made possible by the tractability of this model, reveal a surprising structure of the solution, where a threshold effect is basically at play. Whenever the privacy component of the objective is stronger than freshness, one should completely avoid using the insecure server. Conversely, if average AoI minimization is more important, the insecure server can be used almost indiscriminately to exploit its higher processing rate.**

***Index Terms*—Age of information; Trust; Secure communication; Queueing theory.**

I. Introduction

Age of information (AoI) quantifies the freshness of status reports from a remote source, which is important when tracking real-time content [1]. If the source reports the system status at times $\tau_1, \tau_2, \ldots$, and transmissions $\upsilon_1, \upsilon_2, \ldots$ to reach the destination, so that the reception times are $\tau_1+\upsilon_1, \tau_2+\upsilon_2, \ldots$, then the instantaneous AoI at time t is defined as the time elapsed since the generation of the last received report, i.e., [2]

$$\delta(t) = t - \max_i \{\tau_i : \tau_i + \upsilon_i < t\}\,. \tag{1}$$

This means that $\delta(t)$ is reset at some υ_i values and then grows linearly until a fresher update is received.

Queuing systems are a common scenario for AoI investigations, as analytical results that extend classic evaluations of buffer occupancy or average waiting times are available, often in closed form, allowing precise optimization [3]–[5]. Typically, the instants τ_i are determined as the arrival rate in a queue, and υ_i accounts for the system time spent in the queue, therefore including not only the service time, but also the queueing delay.

We adopt this approach, but focus on a system with a different twist, i.e., including sensitive information and compromised servers. This requires first of all to consider multiple destination servers that can process the information generated by the source. However, this is not a split queue [6], nor a fork-join [7], but rather a redundant queue, which is known to have lower latency (and therefore, in our case, AoI) [8], and for which we derive closed-form analytical expressions.

After that, we consider not only an individual optimization challenge, but also explore the different balance and trade-offs between obtaining a low AoI and preserving privacy [9]–[11]. In particular, we consider a system where the queue can use two servers: one is slower but trusted, the other is faster but potentially compromised. Thus, the source will always use the trusted connection, but can choose to use the other *as well*, which is quantified through a duplication rate ζ expressing how often data are also sent to the potentially compromised server, with the implicit assumption that everything sent to that server can be intercepted by an adversary [12].

We introduce a compound objective, which can be taken as an optimization goal, where the average AoI (to minimize) is combined with the average AoI (to maximize) that an adversary can experience based on the information that is shared to it by the sender. This combination occurs according to a parameter denoted as ω, basically quantifying the importance of not revealing information over lowering AoI.

We obtain interesting results that display some relevant trends. In fact, while determining the optimal ζ would be doable, this would require to solve a complex equation, which can be done through reliable numerical methods, and also may be extremely sensitive to many parameters such as the relative speed of the two servers or the coefficient of importance ω.

It turns out that the final value is crucial for distinguishing between two distinct behavior of the system. In other words, ζ displays a threshold behavior with respect to ω. This qualitative insight into the existence of a threshold policy is a key finding that could be obscured in more complex models that lack closed-form solutions. We identify a critical threshold value, denoted by $\omega_{\rm th}$, which is responsible of determining the behavior of the source, with $\omega_{\rm th}$ typically being near 1. This threshold value for ω directly translates into indicating which one of the two components of the objective is more

This work was supported by the Italian PRIN2022PNRR project "DIGIT4CIRCLE," project code P2022788KK, and by the Federal Ministry of Research, Technology, and Space (BMFTR) in the xG RIC project, part of the research program Communication Systems "Souverän. Digital. Vernetzt." (grant number 16KIS2429K).

important to the source. If $\omega > \omega_{\text{th}}$, the main concern for the source is trusting the security of the server rather than valuing information freshness. In this case, the source does not send any data to the insecure server, which receives $\zeta = 0$. Conversely, when $\omega < \omega_{\text{th}}$, AoI minimization becomes the most important aspect, therefore the duplication rate ζ approaches swiftly 1, nearly disregarding the necessity of employing the insecure server as well [12].

The remainder of this paper is organized as follows. In Sec. II, we discuss the related literature. Sec III presents the system model and derives the closed-form formulas for the scenario, from which one can analyze the performance in terms of the compromise between age and leakage of sensitive information. We display the numerical results in Sec. IV and conclude in Sec. V.

II. Related Work

Classic memoryless queueing systems are a fertile ground for analytical evaluations of AoI, leading to many contributions that typically compute the average AoI for various queue disciplines and service policies [3]–[5], [13], [14]. Most of these papers give a closed-form evaluation that can be used for multiple application scenarios.

In particular, in this paper we consider the formulas related to the simplest LCFS queue with preemption, memoryless arrivals and services, and 1 server, which is denoted in [2] as M/M/1*. We remember that this queue, when taken in isolation, obtains an average AoI Δ equal to

$$\Delta = \frac{1}{\lambda} + \frac{1}{\mu}\,, \tag{2}$$

where λ and μ are the arrival and service rate, respectively. This choice is not restrictive, as the analysis can be extended to other queueing systems, even though sometimes it is at the price of considerably complicating the math [2]. However, in our opinion, an M/M/1* is the choice that makes the most sense if the objective of the analysis is to determine whether the source is willing to use a compromised server to increase the processing speed. This implicitly assumes that the service is oriented towards the fastest service, and thus it preempts older packets with fresh ones. In fact, as argued in Sec. I, we can reduce the v_i to the only service time.

Beyond the queueing theory evaluations that serve as foundation of our analysis, the issues of security and especially confidentiality are relatively less explored in combination to AoI. Most investigations in this context deal with adversaries that actively try to deteriorate the freshness of the exchanged information [15], thus increasing AoI or other related metrics, such as the age of incorrect information. Usually, the adversary is just a jammer [16], [17], and sometimes it can also inject false data to worsen the exchange of fresh information [10], [18]. However, in all these papers, AoI deterioration is the direct effect of the attack, rather than a by-product of information sharing with insecure servers as in our analysis.

A closer similarity to our problem characterizes the papers combining AoI with eavesdropping. In both [9] and [11], part of the information sent to the legitimate receiver can be intercepted by an eavesdropper; thus, the transmitter may decide to decrease the information injection rate. The objective to meet is a multi-variable combination of the age of legitimate information (to minimize) and that of leaked information (to maximize). The difference between these contributions is that [9] considers a queueing system, and eavesdropping happens according to a random probability, whereas [11] focuses on a physical layer analysis. Similar reasoning is also introduced in [12], where a secrecy age metric is introduced to represent the aforementioned trade-off, and a Markov decision process is used to solve the secrecy-age-optimal transmission probability for each state of the system. Thus, in this case, finding the right tradeoff between these contrasting objectives is achieved with a stateful approach, rather than a stationary policy. This is pushed even further in [19] that considers a per-packet scheduling toward AoI minimization but also at risk of eavesdropping, which happens again in random instants with i.i.d. probabilities. Another relevant reference is [20] in which an eavesdropper takes advantage of a timing-based side-channel to gather information from a secure remote pull-based process employing goal-oriented communication.

The combination of AoI and trust is even rarer. In [21], system control is based on trust and AoI indicators, combined together, and fed to a machine learning engine. However, the authors treat these two indicators as orthogonal evaluations for each packet, implying there may be fresh but not very reliable packets, but do not explicitly address how insecure communications may offer a trade-off, improving AoI at the expense of trust.

In [22], a metric called age of trust is proposed, which essentially stems from the same idea as AoI but applied to the instant of the last verification. To the best of our knowledge, these are the only instances of an explicit connection between AoI and trust issues.

Most other references propose a problem for authentication through distributed security mechanisms, exactly to avoid that the information ends up in untrusted destinations [23], [24]. Reference [25] speculates on the inclusion of blockchain validation to this end, but remarks that at the same time this increases the processing time and therefore AoI. This kind of discussion parallels what we consider in the present paper, where a compromised server is faster than a trusted one.

In summary, while prior research addresses issues somewhat akin to ours, they ultimately diverge. On the other hand, there are a couple of studies related to our work in terms of server selection [6], [26], yet our problem remains distinct. This distinction arises because our analysis does not involve selecting a singular server from multiple options. Instead, connections to trusted servers are continuously active and secure, although the source is capable of opting for the compromised server(s) as well to improve the freshness of her data.

III. System model

We consider a scenario where a single source Alice A generates packets according to a Poisson process with rate λ.

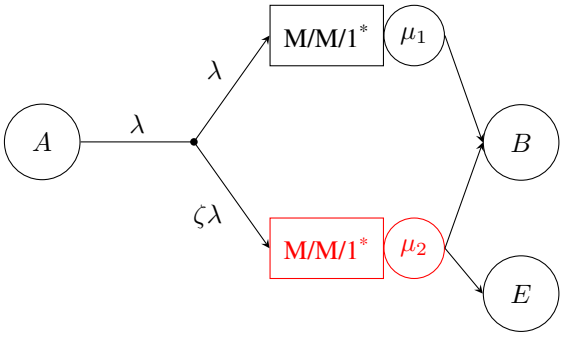

Fig. 1. System model for the dual server scenario.

Alice, as the sender, can decide whether to send traffic only to a secure but slow server or to duplicate it to a fast but insecure server. Both servers are modeled as M/M/1* queues with LCFS processing and preemptive service behavior with exponential service time $Y_1 \sim \text{Exp}(\mu_1)$ and $Y_2 \sim \text{Exp}(\mu_2)$, respectively, with $\mu_2 > \mu_1$. In Sec. III, we also discuss how to extend our results for any $\mu_2 > 0$, independently of its relation to μ_1. Data packets enter the queueing systems with rates λ and $\zeta\lambda$, respectively, with $\zeta \in [0, 1]$ being a duplication rate for the packets. While the trusted server forwards the processed information only to the legitimate recipient Bob B, the insecure server also sends it to a malicious entity Eve E, and this result must be avoided as much as possible by the sender A who is in control of the parameter ζ. Fig. 1 is a graphical representation of the described scenario, where the top server in black is the trusted server, and the bottom server in red is the insecure one. It is important to note that data sent to the servers is the same, and therefore a packet preempted from one of them might still be successfully served by the other.

The average AoI at E is obtained according to [2] as

$$\Delta^E = \frac{1}{\zeta\lambda} + \frac{1}{\mu_2}\,. \tag{3}$$

In contrast, the average AoI for B must take into account the contributions of both queues. Considering the results from renewal reward theory and geometrical considerations on the sawtooth-like shape of AoI evolution, the average AoI expression is [2], [3]

$$\Delta^B = \frac{1}{\lambda} + \frac{\mathbb{E}[Y^2]}{2\mathbb{E}[Y]}\,, \tag{4}$$

where Y is a random variable representing the service time at B. On average, a packet may go to both queues in parallel with probability ζ, and only to the first queue with probability $1-\zeta$. We introduce the random variable $Y_0 = \min\{Y_1, Y_2\}$. Given the properties of exponential random variables $Y_0 \sim \text{Exp}(\mu_1 + \mu_2)$. Therefore Y can be rewritten as

$$Y = \zeta Y_0 + (1-\zeta)Y_1\,. \tag{5}$$

With this expression[1] and the linearity of expectation, it is immediate to compute the first order moment of Y as

$$\begin{aligned}\mathbb{E}[Y] &= \zeta\mathbb{E}[Y_0] + (1-\zeta)\mathbb{E}[Y_1] \\ &= \frac{\zeta}{\mu_1+\mu_2} + \frac{1-\zeta}{\mu_1}\,.\end{aligned} \tag{6}$$

With similar reasoning, the second order moment of Y is obtained as

$$\mathbb{E}[Y^2] = \zeta^2\mathbb{E}[Y_0^2] + (1-\zeta)^2\mathbb{E}[Y_1^2] + 2\zeta(1-\zeta)\mathbb{E}[Y_0Y_1]\,. \tag{7}$$

While the second order moments of Y_0 and Y_1 are known from basic probability theory, the joint moment $\mathbb{E}[Y_0Y_1]$ is not straightforward to obtain as the two random variables are correlated. One way of computing this quantity is by marginalizing the expectation in the two possible values that can be taken by Y_0 and then applying the law of total probability. With this consideration we can write

$$\begin{aligned}\mathbb{E}[Y_0Y_1] = \mathbb{E}\left[Y_1^2 \mid Y_1 < Y_2\right]\text{Prob}[Y_1 < Y_2] + \\ + \mathbb{E}\left[Y_1Y_2 \mid Y_1 > Y_2\right]\text{Prob}[Y_1 > Y_2]\,.\end{aligned} \tag{8}$$

While it is straightforward to compute

$$\text{Prob}[Y_1 < Y_2] = \frac{\mu_1}{\mu_1+\mu_2} \tag{9}$$

$$\text{Prob}[Y_1 > Y_2] = \frac{\mu_2}{\mu_1+\mu_2} \tag{10}$$

thanks to the properties of exponential random variables, we need to do further analysis for the two conditional expectations in (8) because, while Y_1 and Y_2 are independent when taken singularly, they are not independent when conditioned on the value of the other. We therefore need to better characterize the conditional distribution of $Y_1 \mid Y_1 < Y_2$.

Theorem 1. *The conditional random variable $Y_1 \mid Y_1 < Y_2$ is exponentially distributed with parameter $\mu_1 + \mu_2$.*

Proof: Considering the definition of density of a conditional distribution, and noting that the conditioning is applied on a event, we write

$$\begin{aligned}f_{Y_1|Y_1<Y_2}(x) &= \frac{f_{Y_1}(x)\text{Prob}[Y_2 > x]}{\text{Prob}[Y_1 < Y_2]} \\ &= \frac{f_{Y_1}(x)\displaystyle\int_x^{+\infty} f_{Y_2}(y)dy}{\displaystyle\int_0^{+\infty}\int_x^{+\infty} f_{Y_1}(x)f_{Y_2}(y)dydx}\,,\end{aligned} \tag{11}$$

where the numerator is a marginalization over the joint PDF. After some algebra, we obtain

$$f_{Y_1|Y_1<Y_2}(x) = (\mu_1+\mu_2)e^{-(\mu_1+\mu_2)x}\,, \tag{12}$$

[1]A standard mixture model is not sufficient for this analysis. The canonical average AoI formula (4) relies on the assumption that the inter-delivery times are independent and identically distributed draws from a service distribution Y. However, in our system with parallel servers and preemption, this assumption is violated. The distribution of the next successful service time depends on the nature (duplicated or not) of the currently preempted packet. Our approach provides a compact formulation to capture the correlation which would be ignored by a mixture model. The goodness of our model will be validated by the simulations in Sec. IV.

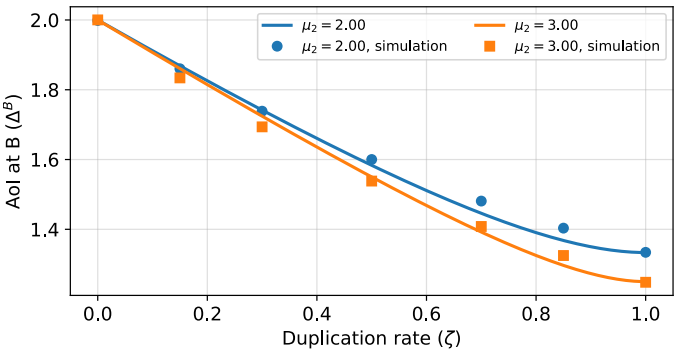

Fig. 2. AoI at B for different values of μ_2. $\lambda = \mu_1 = 1$. Markers indicate simulation results.

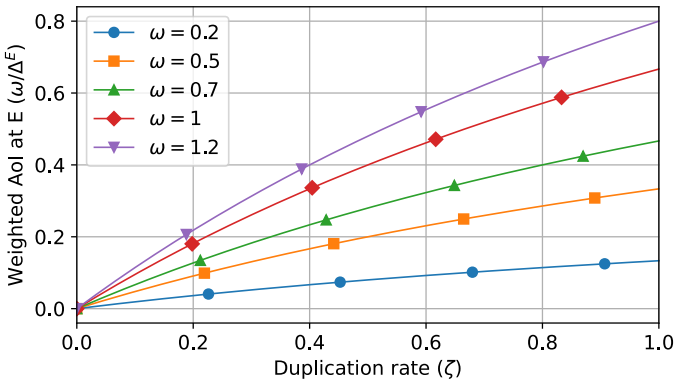

Fig. 3. Weighted AoI at E for different values of ω. $\mu_2 = 2$.

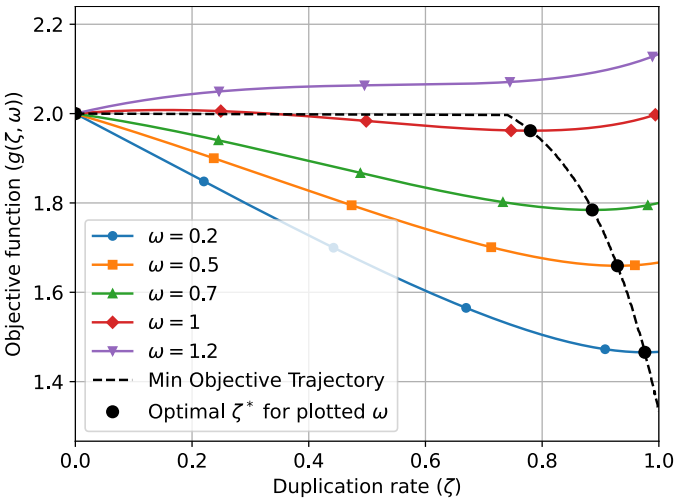

Fig. 4. Objective function for different values of ω. $\mu_2 = 2$.

thus proving the Theorem. ∎

With Theorem 1 at hand, it immediately follows that

$$\mathbb{E}\left[Y_1^2 \mid Y_1 < Y_2\right] = \frac{2}{(\mu_1+\mu_2)^2}\,. \tag{13}$$

The computation of $\mathbb{E}\left[Y_1 Y_2 \mid Y_1 > Y_2\right]$ is a little more involved. It can be noted that, due to the memoryless property of exponential distributions, and given that $Y_1 > Y_2$, we have $Y_1 = Y_2 + Y_1'$, where $Y_1' \sim \text{Exp}(\mu_1)$ and Y_1' is independent from $Y_2 \mid Y_1 > Y_2$. Due to the linearity of the expectation, we can thus write

$$\mathbb{E}\left[Y_1 Y_2 \mid Y_1 > Y_2\right] = \mathbb{E}\left[Y_2^2 \mid Y_1 > Y_2\right] + \mathbb{E}\left[Y_1' Y_2 \mid Y_1 > Y_2\right]\,. \tag{14}$$

To solve this expanded formulation, we can leverage similar reasoning to Theorem 1 and obtain the following result.

Theorem 2. *The conditional random variable $Y_2 \mid Y_1 > Y_2$ is an exponential random variable of parameter $\mu_1 + \mu_2$.*

Proof: The proof closely follows that of Theorem 1 with the due changes in the variables considered. ∎

As a direct consequence of Theorem 2 we promptly obtain a closed form expression for (14)

$$\mathbb{E}\left[Y_1 Y_2 \mid Y_1 > Y_2\right] = \frac{2}{(\mu_1+\mu_2)^2} + \frac{1}{\mu_1(\mu_1+\mu_2)}\,. \tag{15}$$

Combining (13) and (15) into (8) we get the expression of the joint moment of $Y_0 Y_1$

$$\mathbb{E}[Y_0 Y_1] = \frac{2\mu_1+\mu_2}{\mu_1(\mu_1+\mu_2)^2}\,. \tag{16}$$

By including (16) in (7) and plugging the result along with (6) in (4) we finally obtain the expression for the average AoI at the legitimate receiver B served by two parallel M/M/1* servers

$$\Delta^B = \frac{1}{\lambda} + \frac{(\mu_1+\mu_2)^2 - \zeta\mu_2(3\mu_1+2\mu_2) + \zeta^2\mu_2(\mu_1+\mu_2)}{\mu_1(\mu_1+\mu_2)\left[\zeta\mu_1 + (1-\zeta)(\mu_1+\mu_2)\right]}\,. \tag{17}$$

The goal of Alice is to simultaneously balance two contrasting objectives: it wants to reduce as much as possible the AoI at the legitimate receiver Bob while maximizing the AoI for the malicious receiver Eve. To this end, the sender wants to optimize a linear combination of the two AoI measures defined as

$$g(\zeta, \omega) = \Delta^B + \omega\frac{1}{\Delta^E}\,, \tag{18}$$

where we introduce a weighting parameter ω to control the importance for A of the AoI at the illegitimate receiver E. This formulation for the objective function $g(\zeta, \omega)$ is motivated by the fact that Alice is primarily interested in reducing the AoI at Bob's side, while it is willing to allow only partially some information leak to Eve. To this end, we take the reciprocal of Δ^E to avoid the objective to explode when eventually A decides not to send any data to the insecure server.

IV. Results

In this section, we report both analytical solution to the optimization problem and numerical evaluations for the previously defined scenario. In all the following results, we consider $\lambda = \mu_1 = 1$ as this configuration minimizes the AoI for the M/M/1* queue [2].

Fig. 2 shows the AoI on Bob's side as a function of the duplication rate ζ. All curves start from 2 as the only active queue is the one referring to the secure server. As ζ increases, the AoI decreases and eventually stabilizes for $\zeta \to 1$ to reach

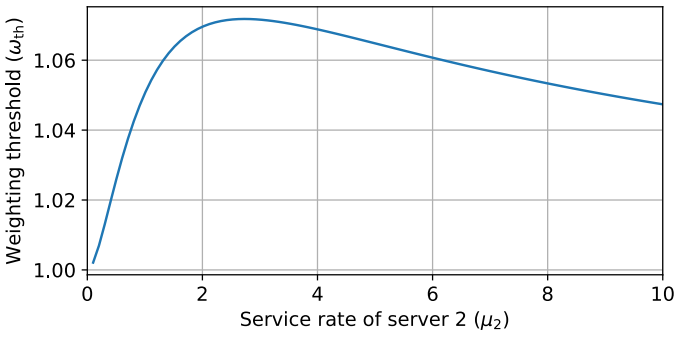

Fig. 5. Threshold value ω_{th} for the separation of the two operations of the system.

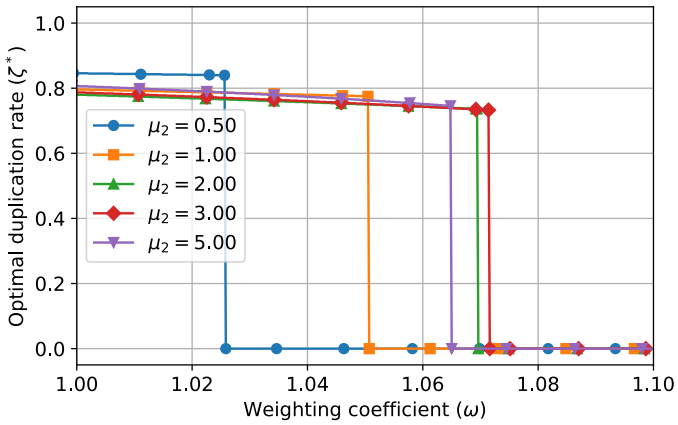

Fig. 6. Optimal duplication rate ζ^* for different values of μ_2.

a minimum. In particular, this improvement in AoI is obtained even for values of $\mu_2 < \mu_1$ since a packet preempted from the trusted server could be served by the insecure one if it was sent as a duplicate to it. We also report the result of Montecarlo simulations run for 10^6 time units to estimate the AoI at Bob's side. The close match between the simulations and our analysis proves the correctness of our analytical solutions. Similarly, Fig. 3 shows the AoI on Eve's side weighted by the parameter ω where $\mu_2 = 2$ is fixed. From the plot, it is evident that higher values of ω increase the concavity of the curve and will eventually increase the impact of this term in the objective function.

Fig. 4 plots the curves of the objective function $g(\zeta, \omega)$ and marks the minimum values obtained for the plotted ω for fixed $\mu_2 = 2$. The ultimate goal for the sender Alice is to minimize this objective function by acting on the choice of parameter ζ. A closed form solution for this problem is quite involved and we will not report the full symbolic expression in the interest of space. However, a precise numerical solution can be easily obtained by numerical means. As ω increases, the optimal value ζ^* that minimizes the objective steadily decreases starting from 1 until a threshold value is reached. At this point, the optimal value for ζ becomes 0, meaning that Alice does not send any more data to the insecure server as this operation is considered undesirable for her objective.

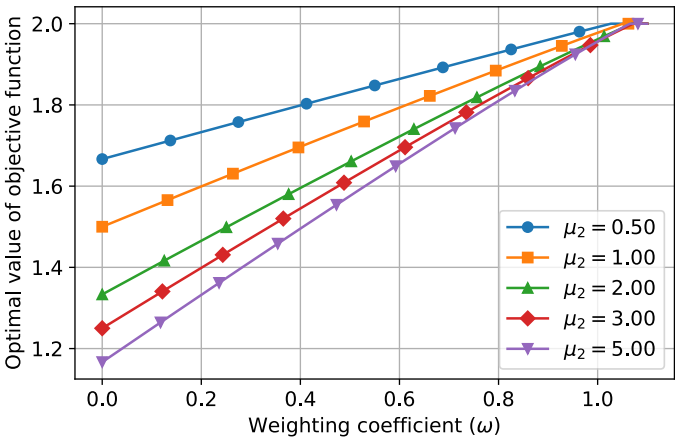

Fig. 7. Optimal value of the objective function $g(\zeta^*, \omega)$ for different values of μ_2.

This result suggests the existence of two distinct operational modes for the system in which A behaves drastically differently. The first working point involves high utilization of the secondary server, and the second one involves not sending any information at all to the auxiliary server. The sudden change in Alice's strategy is controlled by a threshold value for the parameter ω, obtained as the boundary at which the minimization of (18) as a function of ζ has real, non degenerate solutions. Unlike the expression for optimal ζ^*, has an elegant and concise closed-form expression as a function of only the service rate μ_2

$$\omega_{\text{th}} = \frac{2 + 2\mu_2 + 2\mu_2^2 + \mu_2^3}{\mu_2^2 + \mu_2^3} - 2\sqrt{\frac{1 + \mu_2 + \mu_2^2}{\mu_2^4 + \mu_2^5}}. \qquad (19)$$

From this non-trivial expression, we get that $\omega_{\text{th}} > 1$ for $\mu_2 > 0$ and produces a maximum of approximately $\omega_{\text{th}} = 1.07$ for $\mu_2 = 1 + \sqrt{3}$. Interestingly ω_{th} approaches 1 for both the limits $\mu_2 \to 0$ and $\mu_2 \to +\infty$. This result indicates that having an infinitely fast secondary server is not enough for Alice to justify sending more data in an insecure way; moreover, it indicates that there is a diminishing return in Alice's trust if the server is too fast. The complete behavior of ω_{th} is graphically reported in Fig. 5.

Fig. 6 reports the optimal ζ values that minimize the objective function concentrating only on the area close to values of ω_{th}, as for $0 \leq \omega < 1$ the behavior is almost the same as the one shown by the dashed black line in Fig. 4. In this region, we identify the two different regimes in which the system works. For $\omega < \omega_{\text{th}}$, the second server receives almost all the information in duplicate. Interestingly, the optimal duplication rate ζ^* is similar for almost all $\mu_2 \geq \mu_1 = 1$, the only difference being the value of ω at which ω_{th} is reached. The values of $\mu_2 < 1$ have just slightly higher values of ζ^*, but do not indicate a different behavior of the curves. As stated above, when $\omega \geq \omega_{\text{th}}$, Alice decides not to send any traffic to the insecure server regardless of how fast or slow it is compared to the secured one.

Fig. 7 shows the values of the objective function for the optimal values of the duplication rate ζ^* as a function of the weight parameter ω given to the AoI at the illegitimate receiver E. For $\omega < \omega_{\rm th}$, the objective function is strictly increasing, and higher values of μ_2 indicate better performance for small values of ω. This advantage eventually dies out and becomes non-existent for $\omega \geq \omega_{\rm th}$ where the optimal utility is capped at 2 as the only queue in use is the secure one as indicated by Fig. 6.

V. Conclusions

In this work, we have analyzed a communication scenario where a sender is capable of duplicating its traffic to be processed simultaneously by two servers: one secure and one insecure, with the latter potentially offering faster service. We modeled the servers as two simple LCFS queues with preemption, memoryless arrivals and services, and 1 server [2]. The sender aims to optimize a trade-off between two conflicting objectives: minimizing the age of information at a legitimate receiver that collects updates from both servers and maximizing the AoI at an illegitimate receiver that receives updates only to the insecure server. This setup captures the fundamental tension between performance and privacy in systems where information leakage through faster but untrusted channels is a concern.

We derived closed-form expressions for the average AoI experienced by both legitimate and illegitimate receivers under this dual-path strategy. Based on these expressions, we determined optimal configurations for the sender's duplication rate, that is, how frequently updates should be sent to both servers, to strike a balance between timely delivery to the legitimate receiver and limiting the gain in information for the eavesdropper.

Our analysis identified two distinct operational regimes, each dependent on the insecure server's service rate and the sender's tolerance for information leakage. In the first regime, the sender heavily duplicates its updates to both servers. This strategy significantly improves the AoI performance for the legitimate receiver by leveraging the potentially lower latency of the insecure server, albeit at the cost of increased exposure to the illegitimate receiver. In contrast, the second regime corresponds to a conservative approach in which the sender avoids using the insecure server altogether, regardless of its performance, thereby eliminating any leakage risk but possibly sacrificing freshness at the legitimate end.

References

[1] A. Kosta, N. Pappas, V. Angelakis *et al.*, "Age of information: A new concept, metric, and tool," *Found. Trends Netw.*, vol. 12, no. 3, pp. 162–259, 2017.

[2] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1183–1210, 2021.

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

[4] Y. Inoue, H. Masuyama, T. Takine, and T. Tanaka, "A general formula for the stationary distribution of the age of information and its application to single-server queues," *IEEE Trans. Inf. Th.*, vol. 65, no. 12, pp. 8305–8324, 2019.

[5] J. P. Champati, R. R. Avula, T. J. Oechtering, and J. Gross, "Minimum achievable peak age of information under service preemptions and request delay," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1365–1379, 2021.

[6] L. Badia and A. Munari, "Partially stateful server selection for minimal age of information scheduling over a finite horizon," in *Proc. IEEE Infocom ASoI Workshop*, 2025.

[7] F. Chiariotti, B. Soret, and P. Popovski, "Peak age of information distribution bounds for multi-connectivity transmissions," in *Proc. IEEE SPAWC*, 2021.

[8] J. Dean and L. A. Barroso, "The tail at scale software techniques that tolerate latency variability are vital to building responsive large-scale web services," *Commun. ACM*, 2013.

[9] L. Crosara, N. Laurenti, and L. Badia, "Age of information is not just a number: Status updates against an eavesdropping node," *Ad Hoc Networks*, vol. 155, p. 103388, 2024.

[10] P. Kaswan and S. Ulukus, "Susceptibility of age of gossip to timestomping," in *Proc. IEEE Inf. Th. Wkshp (ITW)*, 2022.

[11] L. Zheng, J. Ren, Y. Liu, and Q. Chen, "Analysis and optimization of age of information and age of leaked information for IoT networks," *IEEE Wirel. Commun. Lett.*, vol. 13, no. 12, pp. 3583–3587, 2024.

[12] Q. Wang, H. Chen, P. Mohapatra, and N. Pappas, "Secure status updates under eavesdropping: Age of information-based secrecy metrics," in *Proc. IEEE INFOCOM Wkshps*, 2024.

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

[14] M. Moltafet, M. Leinonen, and M. Codreanu, "Average age of information in a multi-source M/M/1 queueing model with LCFS prioritized packet management," in *Proc. IEEE INFOCOM Wkshps*, 2020.

[15] J. Doncel and M. Assaad, "Optimizing age of information with attacks," in *Proc. NETGCOOP*. Springer, 2024.

[16] G. D. Nguyen, S. Kompella, C. Kam, J. E. Wieselthier, and A. Ephremides, "Impact of hostile interference on information freshness: A game approach," in *Proc. WiOpt*, 2017.

[17] A. Garnaev, W. Zhang, J. Zhong, and R. D. Yates, "Maintaining information freshness under jamming," in *Proc. IEEE Infocom Wkshps*, 2019.

[18] V. Bonagura, S. Panzieri, F. Pascucci, and L. Badia, "Strategic interaction over age of incorrect information for false data injection in cyber-physical systems," *IEEE Trans. Control Netw. Syst.*, vol. 12, no. 1, pp. 872–881, 2025.

[19] F. Yuan, S. Tang, and D. Liu, "AoI-based transmission scheduling for cyber physical systems over fading channel against eavesdropping," *IEEE Internet Things J.*, vol. 11, no. 3, pp. 5455–5466, 2024.

[20] F. Mason, F. Chiariotti, P. Talli, and A. Zanella, "Eavesdropping on goal-oriented communication: Timing attacks and countermeasures," *arXiv preprint arXiv:2411.07088*, 2025.

[21] X. Wang, J. Zhang, C. Chen, J. He, Y. Ma, and X. Guan, "Trust-AoI-aware codesign of scheduling and control for edge-enabled IIoT systems," *IEEE Trans. Ind. Inf.*, vol. 20, no. 2, pp. 2833–2842, 2024.

[22] X. Wang, M. Li, Y. Tao, X. Wang, and H. Wu, "Trust evaluation in mobile crowd sensing networks based on age of trust (AoT)," in *Proc. IEEE TrustCom*, 2024.

[23] A. K. Das, S. Roy, E. Bandara, and S. Shetty, "Securing age-of-information (AoI)-enabled 5G smart warehouse using access control scheme," *IEEE Internet Things J.*, vol. 10, no. 2, pp. 1358–1375, 2023.

[24] L. Giaretta, I. Savvidis, T. Marchioro, Š. Girdzijauskas, G. Pallis, M. D. Dikaiakos, and E. Markatos, "PDS^2: A user-centered decentralized marketplace for privacy preserving data processing," in *Proc. IEEE ICDEW*, 2021.

[25] S. Lee, M. Kim, J. Lee, R.-H. Hsu, and T. Q. Quek, "Is blockchain suitable for data freshness? An age-of-information perspective," *IEEE Network*, vol. 35, no. 2, pp. 96–103, 2021.

[26] Y. Dong, H. Xiao, H. Hu, J. Zhang, Q. Chen, and J. Zhang, "Mean age of information in partial offloading mobile edge computing networks," *arXiv preprint arXiv:2409.16115*, 2024.